

For Reference

NOT TO BE TAKEN FROM THIS ROOM

Ex LIBRIS
UNIVERSITATIS
ALBERTAE NSIS





Digitized by the Internet Archive
in 2023 with funding from
University of Alberta Libraries

https://archive.org/details/Dutton1971_0

THE UNIVERSITY OF ALBERTA

SYSTEM STUDIES USING A SIMULATED COMPUTER

by



J. Barry Dutton

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE
OF MASTER OF SCIENCE

DEPARTMENT OF COMPUTING SCIENCE

EDMONTON, ALBERTA

Spring, 1971

UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies for acceptance, a thesis entitled SYSTEM STUDIES USING A SIMULATED COMPUTER submitted by J. Barry Dutton in partial fulfillment of the requirements for the degree of Master of Science.

ABSTRACT

As the study, at the undergraduate level, of hardware/software systems becomes increasingly important in university curricula, the need for adequate facilities to process the whole range of system oriented assignments becomes critical. Neither the manufacturer's software nor the traditional assembler language processors for student jobs completely fulfill the need. The ideal facility would be to supply each student with an actual computer, complete with software designed to be replaced in stages with student developed software. Since this approach is not economically feasible, a language processor for student programs, the Student Assembler Language Translator (SALT), has been extended to simulate efficiently to a student the characteristics of a simplified /360 computer (SUPERSALT) complete with a replaceable supervisor (SALT MONITOR). How this extended processor might be used is explored in the outline, including assignments to be run on the simulator, of an undergraduate systems course.

ACKNOWLEDGEMENTS

I would like to express my appreciation to my supervisor, Dr. J.P. Penny, for his guidance and encouragement in the preparation of this thesis; to Dr. B. Mailloux for his helpful suggestions and continued interest in the project; and to my wife for her many hours of patient work in typing and proofreading the manuscript.

TABLE OF CONTENTS

	Page
CHAPTER I: INTRODUCTION	
1.1 Undergraduate Education in Computer Systems	1
CHAPTER II: STUDENT PROCESSORS: THE PROBLEM	
2.1 Introduction	4
2.2 The Inconsistent Time Base	4
CHAPTER III: SPECIFICATION OF A SIMULATED COMPUTER SYSTEM	
3.1 Introduction	12
3.2 Specifying the Hardware	12
3.3 Specifying the Software	13
3.4 Performing Input/Output	14
3.5 Replacing the SALT Monitor	19
CHAPTER IV: THE SIMULATED COMPUTER AS AN AID IN PRESENTING A COMPUTER SYSTEMS COURSE	
4.1 Introduction	21
4.2 Using the Manufacturer's Software	21
4.3 Using Special Student Oriented Software	26
4.4 Summary	43

CHAPTER V: IMPLEMENTATION OF SUPERSALT: THE MAJOR PROBLEMS	
5.1 The Internal Facilities Required for CCW I/O Support ...	44
5.2 The Internal Facilities Required to Simulate a Multi- State CPU with Interrupt System	55
5.3 Variations in the System Format	58
CHAPTER VI: LIMITATIONS OF SUPERSALT AND PROPOSALS FOR ITS FUTURE DEVELOPMENT	
6.1 Limitations	60
6.2 Recommendations for Further Development	64
CHAPTER VII: CONCLUSION	72
BIBLIOGRAPHY	74
APPENDIX I: ADDITIONS TO THE MACHINE INSTRUCTION SET	
1 Start I/O - (SIO)	76
2 Test I/O - (TIO)	77
3 Halt I/O - (HIO)	78
4 Set System Mask - (SSM)	79
5 Load Program Status Word (LPSW)	79
APPENDIX II: SUPERSALT HARDWARE CHARACTERISTICS	
1 Specification of I/O Device Characteristics	81
2 Specification of CPU Reserved Memory Format	82

3	Channel Status Word Format	84
4	Program Status Word Format	85

APPENDIX III: THE CHANNEL COMMAND WORD

1	The Assembler Instruction	87
2	The Machine Command Format	88

APPENDIX IV: ADDITIONS TO THE SYSTEM MACRO SET

1	Execute Channel Program	89
2	Test and Wait on I/O	89
3	Channel Command Block	90

APPENDIX V: A SAMPLE SALT PROGRAM	92
---	----

LIST OF FIGURES

	Page
Fig. 2-1 CPU status during uninterrupted execution	5
Fig. 2-2 CPU status does not demonstrate interrupted execution	6
Fig. 2-3 CPU status reflects the synchronization of events occurring within the SUPERSALT computer	9
Fig. 4-1 First stage of the assignment demonstrating conventional subroutine linkage	28
Fig. 4-2 Second stage with user program status words and user program interrupt supervisor	29
Fig. 4-3 Third stage with user CCW's and multiple buffers in place of the READ macro and user supervisor for the WRITE macro SVC	32
Fig. 4-4 Fourth stage with user task and I/O supervisors..	37
Fig. 4-5 Fifth stage employing user assembler for a greatly simplified assembler language	39
Fig. 4-6 Final stage containing user multiple task super- visor running device/memory spool packages asyn- chronously with assembler	41
Fig. 5-1 The device description table	46
Fig. 5-2 The device status table	46
Fig. 5-3 The fields of a pending event to control inform- ation transfer	52
Fig. 5-4 Pending events corresponding to a CCW initiated at time T_1 minus "Average Delay Time"	54
Fig. 6-1 A direct access device addressing scheme	65

CHAPTER I

INTRODUCTION

1.1 Undergraduate Education in Computer Systems

A major objective in the education of computer scientists as opposed to computer users, should be to reduce to a minimum the lack of understanding of computer systems, both software and hardware. Unfortunately it is often true that even skilled users and competent theoretic computer scientists lack comprehension of the principles governing system design. Naturally, such a shortcoming results in an inability to exploit fully the available computing facilities and, potentially more serious, an inability to communicate with those responsible for system design, maintenance and operation. The author has found, for example, that even graduate students, for whom systems studies offer an absorbing field, are often more at ease with the theoretical literature than with the actual systems at their disposal.

During their careers, Computer Science graduates may be expected to engage in such activities as systems programming, consulting, computer management and software engineering, all of which require more than an exclusively theoretical background. They may need to participate in the process of hardware/software selection and thus find it essential to be capable of translating the characteristics of various configurations into terms of operating efficiencies, resource loads and possible bottleneck areas.

Computer Science curricula now appear to be deficient in this

important area of systems studies. An examination of such curricula for universities in Britain [3] and in Canada [1] shows that topics such as assembler language, logic design and language processor construction are usually included. On the other hand, material on computer system structures is frequently absent from curricula and, where taught at all, may be covered only in the final year of the program. Admittedly, the surveys quoted are now three years old, and the situation is certainly changing. The most recent ACM curriculum recommendations [2] include three courses, two intermediate and one advanced, whose synopses show a good coverage of systems topics. However, personal observation suggests that very few universities even now offer such a variety of courses.

Even in those which do, the value of discussing theoretical system concepts is diminished by the lack of opportunity to gain practical experience. Similarly, the teaching of a programming language loses impact if programs cannot be run. The manufacturer's software offers one possible vehicle for gaining a systems background but from the author's experience, it also suffers serious drawbacks. Special student job processors for low level languages are common and might be used as the basis for assignments illustrating system principles, but these have most often been developed to minimize interference with other processing, rather than with a view to providing comprehensive facilities for experimentation. The author was earlier associated with the development of such a system, the Students' Assembler Language Translator (Dutton [7]; Easton and Penny [10]), which is an in-core system for processing student assembler language

jobs on an IBM /360. The original aim was efficiency, but an extended version, designed and implemented as part of this thesis project, has been developed to provide a more complete and realistic simulation of a computer system.

In this thesis we are concerned with:

- (a) exploring the shortcomings and inconsistencies of some current student assembler language processors, particularly SALT, as illustrative tools for systems courses,
- (b) specifying the characteristics of a /360-like pseudo computer that retains the mechanisms fundamental to this genre of hardware while eliminating the non-illustrative detail that makes creating software particularly time consuming and difficult,
- (c) describing some of the main problems encountered in presenting a systems course using the manufacturer's software both as an example for illustration and as the framework within which assignments were done,
- (d) outlining how a similar systems course could have been structured around a project using a comprehensive student system such as that described in (b), had it been available,
- (e) explaining the major problems in extending the SALT system from a language processor with the shortcomings discussed in (a) to a simulator of the computer system specified in (b).

CHAPTER II

STUDENT PROCESSORS: THE PROBLEM

2.1 Introduction

An examination of the student oriented assembler processing systems, for which the author has been able to acquire documentation, reveals that although they provide an excellent simulation of a low level language and certain aspects of the internal machine architecture, they are incomplete or inaccurate in their illustration of several fundamental machine characteristics. The authors of SOS: The Brown University Student Operating System [4] pinpoint a major problem area by indicating that the SOS 'machine' does not overlap its I/O with processing but, on the other hand, they have retained for pedagogical purposes the concept of the data channel. It is not clear what pedagogical purpose is served by a channel incapable of asynchronous operation with the CPU, since its existence thus seems pointless.

The other major problem area is the absence of an explicitly defined interrupt mechanism in the simulated processors. This mechanism is completely absent in all three student systems: SOS, STASS and SALT. The following discussion will explore in more depth the problems as they occur in SALT, with the understanding that a similar discussion could be made of either SOS or STASS.

2.2 The Inconsistent Time Base

Consider the following program segment:


```

REPEAT  LA    R3,5(R3)
        LA    5,10(5)
        MVC   0(20,3),0(5)
        CR    5,3
        BH    REPEAT
        .
        .
        .

```

No inconsistencies, for the purpose of illustration, result from the execution of these instructions on the hypothetical SALT computer. The memory organization and addressing scheme are faithfully reflected and the CPU procedure for calculating an address is easily described. The existence and nature of micro-programs can be readily illustrated with reference to the machine instruction formats. The existence of general purpose registers is apparent, and special purpose registers and conditional branching control to explain in detail the inner workings of a central processor, are easily postulated. Most important, the time base is realistic since during interpretive execution of the instructions, accurate execution times (for a /360 Model 65) are calculated by the interpreter for each instruction (Fig. 2-1).

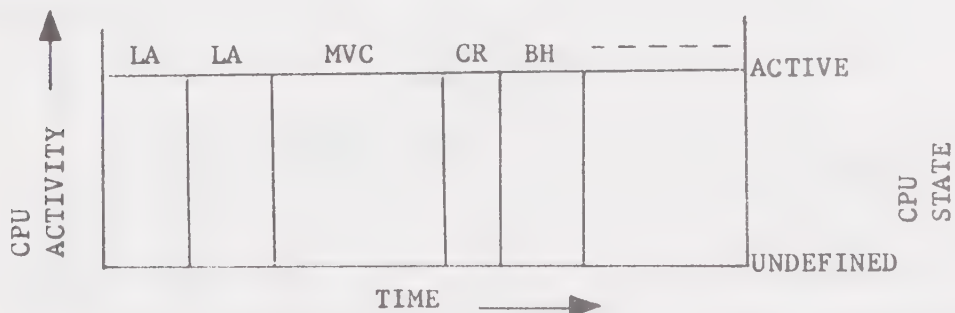


Fig. 2-1. CPU status during uninterrupted execution

However, consider a program segment containing an I/O request:

```

                LA    5,10(5)
                READ  CARD
+              LA    1,CARD
+              SVC   1
                MVC   0(4,5),CARD+20
                .
                .
                .
CARD          DS    80C

```

The I/O request is realistically interfaced to the SALT hardware through a system macro "READ" which places the input buffer address in register 1 and generates a supervisor call to the SALT monitor.

The activity performed by the monitor is, however, impossible to describe precisely since the instruction following the supervisor call appears to be executed without an intervening pause and may, in fact, refer to the information just obtained in the READ operation. That is, the record inexplicably appears in memory in the instant between the execution of the SVC and the following instruction. This discontinuity in the time base is illustrated in Fig. 2-2.

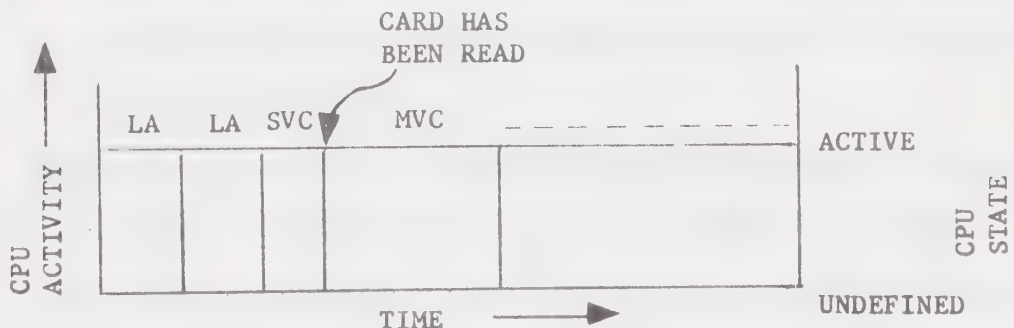


Fig. 2-2. CPU status does not demonstrate interrupted execution

Besides the inconsistent time base, SALT has several other shortcomings which severely limit its usefulness for illustrating the principles of hardware/software systems. These are:

- (a) Communication between the CPU, in which SALT instructions appear to be executed, and its environment is not well defined to the user since I/O devices are connected, and operate in an invisible fashion.
- (b) The characteristics that can be ascribed to the SALT CPU from examination of SALT machine instruction execution in SALT time indicate that the CPU has a single state, i.e., running.
- (c) Even if other hypothetical states are postulated for the CPU, there is no hardware mechanism, defined and visible to the student, for changing the CPU state, for terminating or switching the execution of instructions from a given instruction stream, or for initiating CPU operation at some time T.
- (d) Even though the existence of software support for the student's program is demonstrable, and the existence of hardware is obvious, it is not at all clear precisely where hardware ends and software begins or how the one is interfaced to the other.

Appendices I to IV list detailed specifications for a minimal, /360-like machine, SUPERSALT, with software support such that:

- (a) Pseudo programmable channels with pseudo devices attached are visible to the user, and these channels appear to oper-

ate asynchronously with the SUPERSALT CPU in a fashion consistent with the simulated time base.

- (b) A multi-state CPU with precisely defined hardware characteristics is apparent and available to the programmer, either as a bare machine on which user-supplied software can be employed to run user programs, or in combination with the SALT monitor employed to run user-supplied programs, or in virtually any combination of SALT monitor and user software support.

The following program segment, using the facilities described in Appendices III and IV, would, when executed on SUPERSALT under the SALT monitor, appear to the user as having a fully consistent time base (Fig. 2-3).

```

LA      5,10(5)
EXCP    CRDRDR
+       L      1,=A(CRDRDR)
+       SVC     6
MVC     0(4,5),0(3)
LA      3,10(3)
WAIT    CRDRDR
+       L      1,=A(CRDRDR)
+       TM      10(1),X'01'
+       BO      *+6
+       SVC     7
L       5,0(3)
.
.
.
```

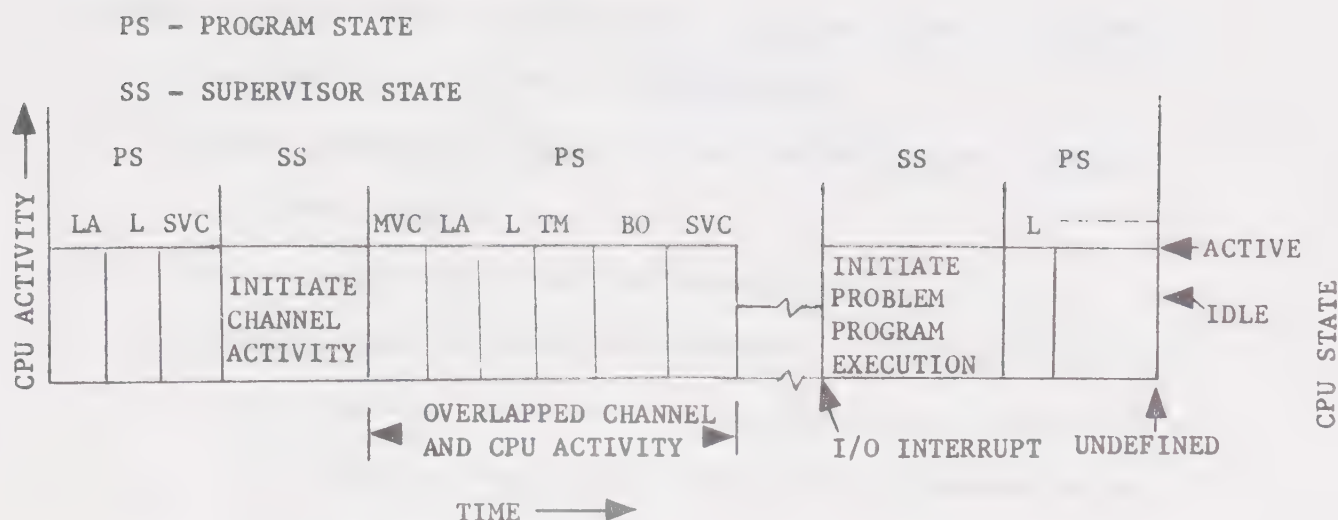



Fig. 2-3. CPU status reflects the synchronization of events occurring within the SUPERSALT computer

Before specifying the characteristics of a student-oriented processor suitable for running students' problems in a systems course, it is first necessary to make some assumptions concerning which aspects of hardware/software systems might be required for student use. The author believes that the most useful facility for illustrating systems concepts would be a set of hardware (simulated or real), complete with software which the student could replace with his own. Since actual hardware is not normally available for the exclusive use of individual students, and since the already existing SALT system offered an excellent base on which to build a simulator, the author decided to proceed from that point. This decision restricted the simulated hardware to the possession of characteristics somewhat like those of an IBM /360 computer. However, in those instances where the author had the opportunity to influence the characteristics of the hardware/software combin-

ation, he did so with the intention of creating a tool of suitable power and complexity for use by students who:

- (a) had completed at least one, and preferably two, Computing Science courses,
- (b) had mastered an assembler language, and
- (c) had reached the level of development at which they had become somewhat blasé about programming algorithms, and were interested in learning more about the environment in which their programs were run.

They could then choose to take as their first course in computer systems, one in which the following areas are explored:

(a) CPU characteristics:

- (i) the nature of instruction execution by the CPU,
- (ii) the need for multiple CPU states to control instruction execution and the instruction set recognized by the CPU at any given time,
- (iii) the concept of instantaneous CPU status,
- (iv) the need for a mechanism, both to save the instantaneous CPU status and to reset it to new, predetermined values in order to change CPU states, and to switch the instruction stream being executed.

(b) Channel characteristics and channel/CPU interaction:

- (i) the nature of a channel as a primitive processor and the reasons for making it programmable,
- (ii) the concept of multi-processing as exemplified by asynchronous, overlapped operation of channel and CPU,

- (iii) the need for making the channel capable of interrupting CPU instruction execution to facilitate synchronization and control of I/O and program operation.

(c) Software:

- (i) software as the interface between the hardware and the program,
- (ii) the concept of continuous, protected operation through software management,
- (iii) the use of software to resolve speed mismatches between hardware components ,
- (iv) software systems for efficient machine utilization: the concept of "task" and the synchronization constraints between tasks; multiprogramming within job bounds and outside job bounds,
- (v) introduction to data management: external storage; file organization,
- (vi) storage management and memory protection,
- (vii) the effects of device characteristics and system configurations on system tuning,
- (viii) initiating processing on the hardware/software combination.

CHAPTER III

SPECIFICATION OF A SIMULATED COMPUTER SYSTEM

3.1 Introduction

An extended SALT system, simulating a computer called SUPERSALT, has been designed to provide students with facilities normally present in an actual hardware/software system, but not usually available or visible to the programmer. The student is given the facility to specify within each job to be run on SUPERSALT:

- (a) a configuration of peripherals and their characteristics,
- (b) whether the SALT monitor is to reside in SUPERSALT or whether the job contains the user's own monitor to supervise the execution of his program.

As well as specifying the hardware configuration, the user may program the peripheral devices from his program and have the execution of these channel programs overlap the execution of his SALT instruction. If the user elects to include, as his software configuration, his own monitor, then he essentially has complete control of the simulated hardware including the handling and masking of interrupts and the testing and initiation of channel operation.

The SALT system, however, maintains control over the transition between jobs, and prints out statistical and diagnostic information as required.

3.2 Specifying the Hardware

The peripheral configuration is specified on the \$SALT card in

the format shown in Appendix II. Currently the peripherals are restricted to pre-set, fixed length records, read or written in a strict physical sequence that is not re-transmittable, i.e., a record once read cannot be re-read. In these characteristics, the peripherals are similar to card readers or line printers, but they differ in that the average rate of delay before transmission (the time required to physically read or write the record at the device) and the data transmission rate (the rate at which the data is transmitted across the channel) are specified by the programmer. These reader/printer-like devices and their channels can, through changing their characteristics, be made to operate at rates varying from typical unit record values up to main memory speed, thus demonstrating CPU usage as a function of CPU/peripheral speed mismatch. The characteristics of a channel other than the channel address or data rate are pre-set. A channel can have more than one device attached to it but it can service (execute a channel program for) only one device at a time. All channels have data widths of 160 bits, i.e., they are capable of transmitting 20 bytes in parallel. These 160 bit wide transmissions follow one another serially at intervals of $(\text{channel width} \times \text{constant}) / (\text{data transfer rate})$ seconds.

3.3 Specifying the Software

The software configuration also is specified on the \$SALT card through the presence or absence of the relative address of a reserved memory location in the program into which the SUPERSALT hardware can exchange status information. If the relative address is absent or

invalid, then the status exchange mechanism is not visible to the programmer, and all requests for supervisory functions are serviced by the SALT monitor. If the relative address is present in the format specified in Appendix II, then the hardware functions of status exchange, initiation of channel operation, timer setting and memory protection bound definition are placed under user control through the values in the specified reserved core area. The format of this reserved memory is shown in Appendix II. Programmer control of these hardware functions normally implies that a user monitor forming one part of a SALT job will service supervisor requests from a problem program which forms another part of the same job.

The extended system, however, maintains overall control over individual jobs by monitoring execution time and pages of output, and overseeing job-to-job transition. The system may be called upon to perform supervisory functions which the programmer chose not to or could not accommodate in his own monitor.

3.4 Performing Input/Output

If the "DEVICE=" option was specified on the \$SALT card then the user may include, as part of his SALT job, channel programs for the channels connecting the devices to the CPU, and he may use SALT system macros to:

- (a) cause initiation of individual channel programs,
- (b) relate channel programs to a given channel and device,
- (c) impose those synchronization constraints between channel

operation and instruction execution that are required by the logic of his program.

The characteristics of the channel command words (CCW's) that make up channel programs are presented in Appendix III and the system macros used to control channel execution are described in Appendix IV.

SUPERSALT channels, once activated, behave as elementary computer processors. They execute, in place of instructions, CCW's retrieved sequentially from main memory by the channel. By executing a special instruction within the CPU a channel is directed to fetch the first CCW of a channel program (a set of logically connected CCW's). However, after activation, the channel appears to fetch and execute CCW's asynchronously with the CPU execution of instructions. Of the four CCW commands, only two (READ and WRITE) can cause data to be transferred across the channel. Since currently supported SUPERSALT devices are considered as having their own internal buffer storage, the reading or writing of records is not concurrent with the transmission of data across the channel. The execution of a CCW READ command causes the record information to be inserted into the reader's internal buffer from where it is transferred down the channel in amounts determined by the channel program. Similarly, on output, information is transferred under control of a channel program from main memory to part or all of the buffer storage in the output device, after which the entire contents of the buffer become the output record. Thus a single CCW cannot cause more than one record to be written or read, but a

single record may be written or read by a number of CCW's chained together (see Chain Data in Appendix III). Similarly, another type of chaining may cause several records to be read or written by a string of CCW's forming a single channel program requiring only one initiation by the CPU (see Chain Command in Appendix III).

Two other commands can be issued in a CCW. The CONTROL command causes mechanical activity such as spacing or skipping on the printer. The TRANSFER IN CHANNEL command causes no activity at a device, but instead, it instructs the channel to fetch a new CCW, not the one in the next contiguous eight bytes as would happen with chained CCW's, but rather the eight bytes starting at the address specified in the "addr" field of the TRANSFER CCW. It is analogous to an unconditional branch instruction.

A great deal of care was taken to ensure that I/O activity appears to take place along a realistic but uncluttered time base. A typical sequence of events for a READ operation might be:

- (a) CPU execution of a load instruction (part of the expansion of an EXCP macro) places the address of a CCB in register 1.
- (b) CPU execution of the EXCP SVC instruction results in a SUPERSALT SVC interrupt.
- (c) If the user has provided his own supervisor as is discussed in Section 3.5, the next event is the CPU execution of the instruction at the address specified in the new SVC PSW,
- (d) If no user supervisor has been provided (which is the situation assumed in (e) to (h) below), the SALT monitor services

the SVC by initiating, if possible, the execution of the channel program.

- (e) At a SALT time of 100 μ sec after the execution of the SVC, control returns to the user-supplied instruction which follows the EXCP SVC. (This is the time required by the supervisor to initiate I/O.)
- (f) Instruction execution continues until either a WAIT SVC is executed or a time equal to the specified average delay time for the device has passed. In the latter case twenty bytes of the input record under CCW control are placed into memory.
- (g) Instruction execution continues with subsequent twenty byte portions of the input record being stored according to CCW specifications at intervals of $(20/\text{data transmission rate}) \times (\text{constant}) \mu\text{sec}$. This continues until either the channel program terminates or a WAIT SVC is executed.
- (h) If a WAIT SVC is executed, the CPU appears to be placed in the WAIT state for a length of time equal to that required to complete the channel program, after which control returns to the instruction following the WAIT SVC.

The system macros of Appendix IV will not execute properly unless the "DEVICE=" option has been entered on the \$SALT card, and the devices and channels have the same addresses as those specified in the Channel Control Blocks. More than one device may be attached to a single channel but, since a channel is wholly dedicated to a device during the time when a channel program is being executed, any

attempt to perform I/O on the other device will cause abnormal termination of the job (if run under the SALT monitor). That is, it is the programmer's responsibility while using these system macros under the SALT monitor to synchronize his I/O requests in order to avoid resource contention, since the SALT monitor does not queue multiple requests for the same resource. It is, however, quite possible to have separate channels controlling separate devices, operating simultaneously with CPU execution of SALT instructions.

If the "DEVICE=" option is not present on the \$SALT card, the sole means of performing I/O becomes the original READ/WRITE system macros described by Penny and Easton [10]. Since no channel or device address is associated with these macros, they must be considered as performing I/O on the "system input" or "system output" devices.

If the "DEVICE=" option is present, the READ/WRITE macros can still be executed; however, no CPU channel overlap results. For example, the execution of a READ would cause a record to be read on the unit specified in the "DEVICE=READER,..." field (if one is present) on the \$SALT card. Control would be returned to the instruction following the READ SVC only after the entire record had been transmitted to the address specified in the READ. When a READ or WRITE is executed, care must be taken to ensure that no channel program, initiated by an EXCP macro, is in execution on the required device since the resulting I/O will be a combination of the records. The READ and WRITE can be considered a higher level of logical I/O support than the EXCP.

Examples of channel programs and the corresponding system

macros appear in Appendix V. A detailed description of channel programming appears in Dutton [5].

3.5 Replacing the SALT Monitor

If the "PSW=" option appears on the \$SALT statement the SALT program, once assembled, is treated as a stand-alone program running on a "bare" machine. Although system software support is still present to take over in case the SALT program cannot maintain execution on the bare machine or in case convenience services are requested by the executing program, this underlying level of software is kept invisible as much as possible. The SUPERSALT machine is defined to be consistent and self-contained within the limits of execution of one SALT job. That is, it has:

- (a) a visible interrupt system to which the user must interface his SALT program,
- (b) a program activated supervisor state in which the CPU recognizes special instructions,
- (c) a CPU that is either executing (or prevented from executing) along a consistent time base,
- (d) provision for a memory protection mechanism (not currently implemented),
- (e) provision for a timer to be decremented in SALT time (not currently implemented).

The exact relationship between the user software on the "bare" machine and the SALT monitor is that any program or SVC interrupt occurring when the SUPERSALT CPU is in the supervisor state is

intercepted by the monitor without reference to the user PSW's. Conceptually this is undesirable since the SALT monitor appears to run on an undefined machine, but practically it is both a necessity to maintain the flow of SALT jobs (by placing limits on execution time, lines printed and level of interrupts) and a considerable convenience to the user since he can use the facilities of the monitor during development of his own software by reflecting selected interrupts to it. In any case, since there are no instructions in SUPERSALT that operate on packed decimal fields, the process of converting from EBCDIC to binary and vice versa should be reflected to the monitor, and since job-to-job transition is beyond the user's control, the normal end of SALT job must be accomplished through the monitor.

An example program containing user PSW's and a simple set of interrupt handlers can be found in Appendix V. A more detailed description of interrupts and PSW's appears in Dutton [5].

CHAPTER IV

THE SIMULATED COMPUTER AS AN AID IN PRESENTING A COMPUTER SYSTEMS COURSE

4.1 Introduction

The system specified in the previous chapter might be used to advantage when setting assignments in a course dealing with software/hardware systems. The course would most naturally follow those in which the student had been exposed to high level languages and to some Assembler Language, and might form a third course in Computing Science. The facilities offered by the extension appear to provide a solution to the problems encountered by the author in setting assignments for such a course, and also appear to overcome the shortcomings, for purposes of instruction, of the manufacturer's hardware/software system. The following section, 4.2, describes some of these problems experienced by the author when teaching a third-year course in systems, using only the manufacturer's software. Section 4.3 outlines a hypothetical, long term, modular project using SUPERSALT as the basis for the assignments in a somewhat similar course.

4.2 Using the Manufacturer's Software

During the first term of the 1970-1971 academic year at the University of Alberta, the author was responsible for organizing and presenting a third-year course consisting of the study of a particular operating system - in this case, the IBM /360 Operating

System. The approach taken was to study the system from three different points of view:

- (a) that of the applications programmer,
- (b) that of the systems programmer, and
- (c) that of the system designer.

Before examining the system structure in point (c), the class was presented with a description of the characteristics of the /360 CPU and the relationship between channel and CPU operation. It was felt that no explicit understanding of the characteristics and operation of the software could be possible without a basic knowledge of the hardware to which the system was designed to interface.

A set of five programming assignments using the manufacturer's software under /360 OS-MVT was designed to provide, as far as possible, practical illustrations of the more important aspects of the lecture material. These assignments were constructed in such a way as to build upon each other and to emphasize in succession:

- (a) the system command language (JCL),
- (b) the Assembler Language features,
- (c) the data set organization and several levels of data management facilities,
- (d) the system utilities and external storage organization,
- (e) the system macros,
- (f) the dynamic program structure and multi-tasking.

An examination of these points reveals that the first three illustrate some aspects of the operating system from the programmer's point of view, and probably all six are of concern to the systems

programmer; however none of them provides a detailed insight into hardware interactions, and none can be considered particularly illustrative of the internal structure of the software. Perhaps the last assignment provided a slight indication of systems operation in that it consisted of a main program that ATTACH'ed three asynchronous subtasks, one to read the input data, while another processed the card image and a third printed out the results. This organization slightly resembles the spooling function required in a multi-programming environment. In fact, during the latter part of the course it became necessary to explain, without student opportunity for practical observation, the other side of the picture, i.e., what was required in the way of system software to "accommodate" the students' programs on the hardware. It does not appear feasible to illustrate through programming assignments using the /360 OS-MVT software, this most important aspect of systems study (short of allowing each student access to the actual hardware).

Another problem that became quite perplexing toward the end of the course was how to determine the level of detail to present in the lectures. During the first part of the course which dealt with programming and systems programming, it was quite feasible and desirable to conduct the lectures on a general plane. The student was expected to fill in the details through massive reading assignments using the manufacturer's manuals. However, during the latter part of the course when the emphasis shifted toward studying the internals of the system, the lectures became almost the sole source of information since the manufacturer's Operating System Program Logic Manuals

were hopelessly detailed, and yet they formed the only comprehensive description of the system structure that the author had encountered. This made it difficult to attain an appropriate level of complexity. The following approaches were considered:

- (a) formulation of the system as an idealized abstraction,
- (b) presentation of the system by describing its control blocks, queues, etc.,
- (c) presentation of a combination of the above, a conceptual description supported by a somewhat simplified description of the main control blocks and other system components.

The first alternative was rejected since it contravened the original aim of the course, to provide a pragmatic view of an actual operating system. The second was not seriously considered since there was insufficient time to present even a small fraction of the bulk of detail involved, and much of it would have been irrelevant to a "functional" understanding of the system in any case. The third alternative was chosen and found to be feasible, although far from ideal. The author found it disconcertingly easy to eliminate control blocks and other components thought to be superfluous to the understanding of the functioning system, only to find they (or some part of them) were essential in explaining some other characteristic of the software. Also, when a system function was reduced to its essentials, it must have seemed confusing to the student who felt he understood the function, to later encounter in his outside reading other seemingly unnecessary components.

A third problem, one that appears to plague any course which

uses the more exotic facilities of a software system, was the high risk of causing malfunction of the resident software through student programming errors, thus interrupting normal service to all users. No amount of precaution can completely eliminate this risk. In an attempt to pinpoint sensitive areas, all five of the assignments for the course were programmed before the course began. When each assignment was handed out, a special workshop session was held, in which the original programmer of the problem discussed these areas and suggested how the problem might be programmed. In addition, each student program was run under a small monitor in an attempt to isolate that program from the system to some small degree.

In spite of these precautions, a few systems failures were traced definitely to the students' programs, and several more failures were suspected to have been caused by them. In any event, some disturbance to normal system operation was experienced, although not to such a degree as to cause revision of the assignment plans.

A final factor that must be considered for any course in which programming forms the bulk of the assignments is the cost of resources required from the computer installation for the debugging and running of student programs. In total, the five assignments mentioned above required approximately:

- (a) 10 cylinders of 2314 disk space reserved over a period of two months,
- (b) 137 computer runs (separate OS jobs, each requiring 100 K of memory) per student,
- (c) 80 minutes execution time on a /360 Model 65 per student,

for 23 students. It should be emphasized that the figures in (b) and (c) above were obtained by extrapolation of incomplete data. However, even if these figures are taken only as a very rough guide, they indicate a considerable investment of resources for a single half-course.

4.3 Using Special Student Oriented Software

The SALT system described in Chapter III could be used to support a set of assignments around which a system course, similar to that discussed in the previous section but without the attendant problems, could be built. Each assignment is designed to build on those preceding it as did the set described in 4.2, but the intention in this case is to develop in parallel both a user program of progressively greater complexity and the supervisor functions necessary to support it. In its final form, the student will have created a full supervisor under which runs an assembler of his own creation, which uses supervisor functions to assemble and load programs which in turn use supervisor functions during their execution. The lecture material which these assignments are designed to illustrate would cover much the same ground as the course described in Section 4.2 but the manner of exposition would be radically different. In the first case, an attempt was made to simplify (and generalize from) a currently existing operating system of great complexity as an example of a typical operating system. In the other case, the lectures would attempt, given the SUPERSALT hardware characteristics, to explore a variety of software configurations as solutions to varying

design aims.

The first assignment (Fig. 4-1) is intended to demonstrate modularity in a SALT program and the use of standard linkage conventions such as those used in OS, as the interface mechanism between modules. The assignment uses standard SALT system macros and the familiar SALT monitor to provide supervisor facilities. The logic of the modules need not be extensive since the important factor in the assignment is the overall structure.

While the assignment is in progress the lectures could most profitably be centered on a description of the SUPERSALT hardware under the following topics:

- (a) machine instruction format,
- (b) instruction fetching and separation into component fields,
- (c) address generation,
- (d) instruction address register (IAR),
- (e) memory address registers (MAR),
- (f) current CPU status,
- (g) the interrupt - an exchange of status information,
- (h) the program status word (PSW),
- (i) memory protection mechanism,
- (j) problem/supervisor state,
- (k) wait/run state,
- (l) interrupt classes.

At this point, an introduction to software might include a discussion of the relationship between assembler language and machine instructions.

Assignment #2 (Fig. 4-2) is intended to make use of the lecture

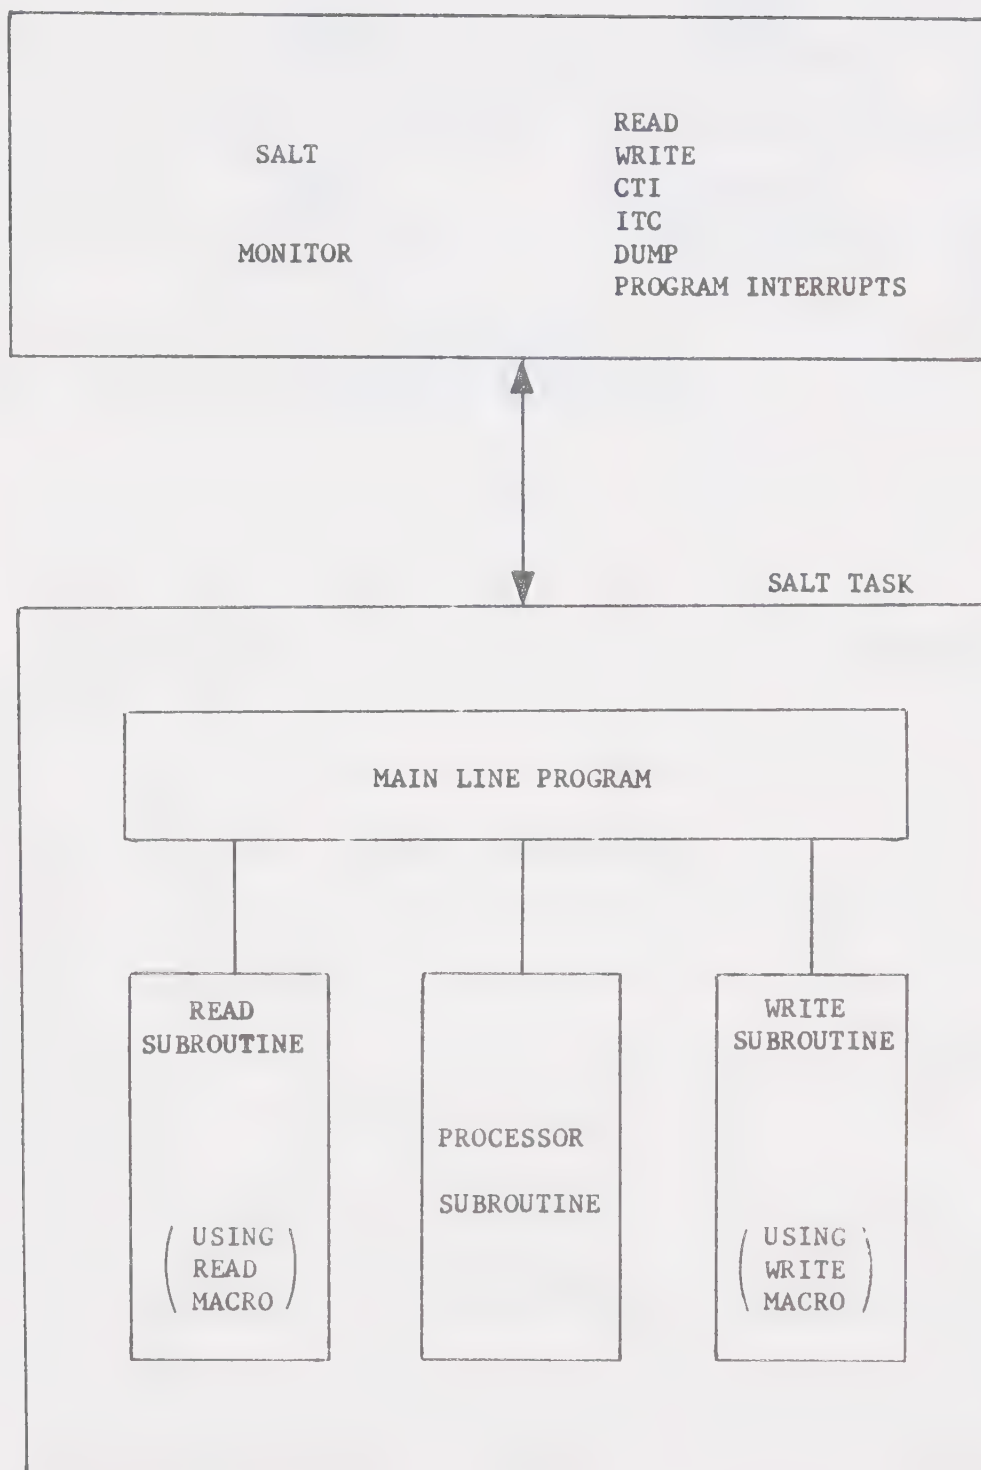


Fig. 4-1. First stage of the assignment demonstrating conventional subroutine linkage

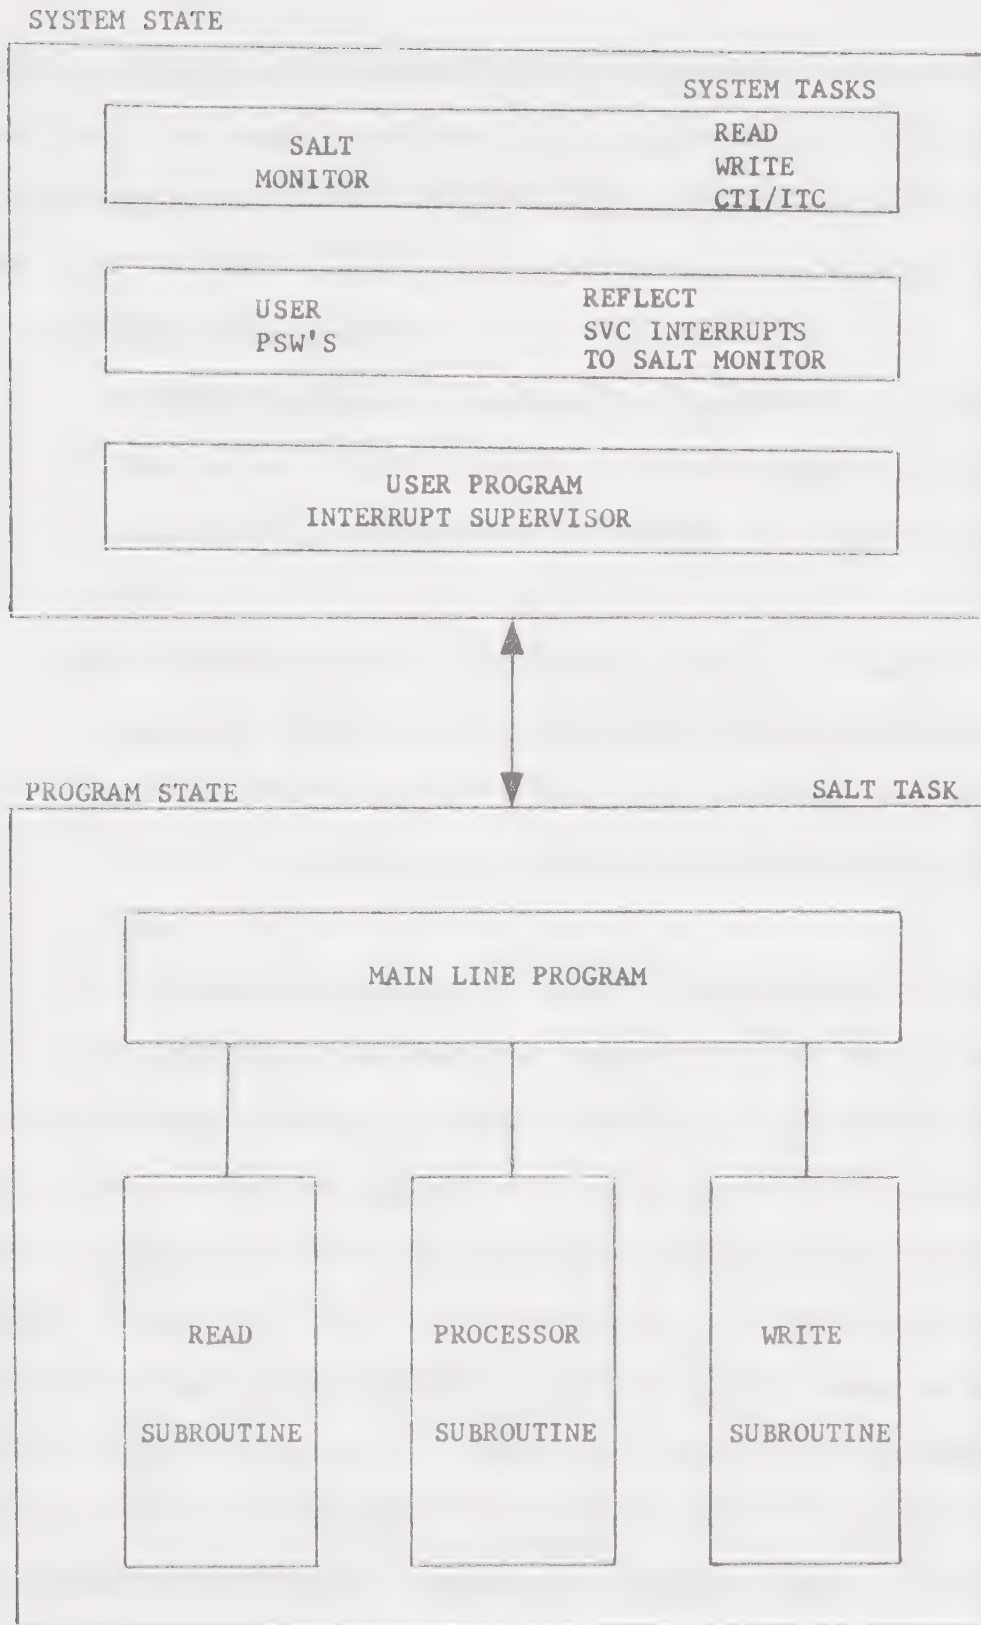


Fig. 4-2. Second stage with user program status words and user program interrupt supervisor

material on SUPERSALT hardware in that the student is required to include, with his program from the first assignment, a set of program status words at the address specified in the "PSW=" option on the \$SALT card. The new PSW's specify user interrupt handling routines with the following characteristics:

- (a) For SVC interrupts: a routine that executes, in the supervisor state, an SVC instruction with the same SVC code as the original, thus passing the request on to the SALT monitor.
- (b) For I/O interrupts: a routine that issues an "impossible condition" message since only READ and WRITE macros have been used and their SVC's have to be reflected to the SALT monitor; thus no I/O interrupts should be visible to the user.
- (c) For program interrupts: a routine that performs, for the user, the debugging function of a program interrupt supervisor.

The program interrupt supervisor can be as comprehensive as the user wishes to make it. Typically it would issue error messages and dump the register contents and selected portions of memory just as the SALT monitor would have done. The DUMP system macro would be a convenient means of accomplishing this since its SVC would be reflected to the SALT monitor. A more sophisticated version might attempt to continue execution by terminating the offending subroutine, reloading the caller's registers by following back the save area links and then continuing processing to the next interrupt. In any case an inventive user should be able to decrease the debugging

time for subsequent assignments by obtaining more information from each run.

The lecture material, while assignment #2 is being completed, would center on the following topics:

- (a) channel operation and device characteristics,
- (b) the channel command word and channel programs,
- (c) the system macros EXCP, WAIT, CCB,
- (d) overlapped I/O and buffer management,
- (e) logical and physical level I/O.

Assignment #3 (Fig. 4-3) is intended to use this information concerning channels and devices to demonstrate what is involved in supporting a logical level of I/O for the WRITE subroutine of assignment #2 and to replace the user's READ subroutine with one containing channel programs.

In assignment #2 the major change to the user's supervisor is the modification of the SVC primary interrupt handling routine to recognize and intercept WRITE SVC's rather than passing them on to the SALT monitor. Interception of the WRITE request involves supplying the supervisor routine to support the following characteristics of the logical I/O:

- (a) The WRITE macro specifies no device address and is assumed to request output on the system output device of a 120 byte record with carriage control.
- (b) When control returns to the instruction following the SVC, the record is assumed to have been written since there is no overlap between logical I/O and instructions in the

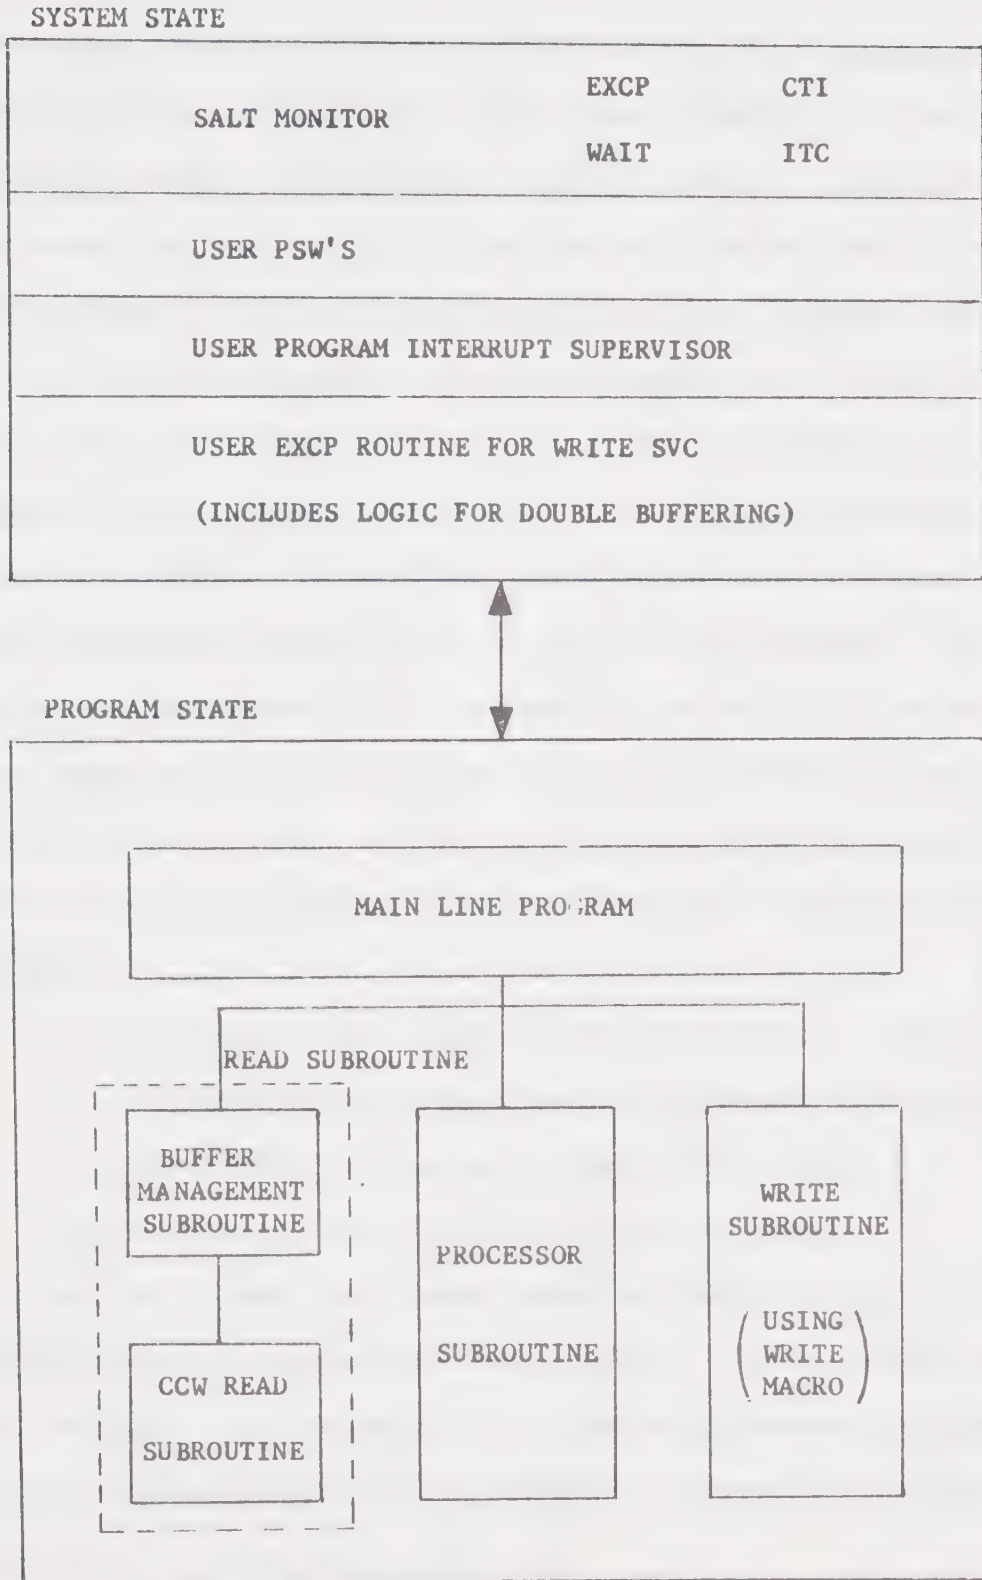


Fig. 4-3. Third stage with user CCW's and multiple buffers in place of the READ macro and user supervisor routine for the WRITE macro SVC

issuing program.

The support routine would use the EXCP, CCB and WAIT group of macros to control the execution of the channel program equivalent to the WRITE request. The EXCP and WAIT SVC's would be issued in the supervisor state and therefore they would be serviced by the SALT monitor. In order to provide some overlap of channel and CPU activity, the support routine would manage two internal buffers filled from the program I/O area by MVC's and emptied by channel programs, employing the strategy of keeping them empty whenever possible. The WAIT SVC should be executed only when a WRITE request has been made and both buffers are occupied, one buffer in the process of being emptied by the currently executing channel program and the other waiting to be emptied as soon as the channel is free. The WAIT, of course, would be issued on the currently active CCB so that when execution resumed following the WAIT:

- (a) the other channel program could be initiated by EXCP,
- (b) the current WRITE request could be serviced by moving the output record to the newly freed buffers, and
- (c) execution of the problem program could resume.

The user's supervisor should intercept the problem program EOJ SVC in order to issue WAIT macros on each CCB in the WRITE support routine. This ensures that all the output records have been written from the buffers before the SVC is reissued in the supervisor state to cause end of SALT job.

While the WRITE support routine as part of the software system

must be designed to supply a general function, making no assumptions on the characteristics of the programs it is to service, the user's new READ routine, as part of a problem program, can be tailored to the needs of that program. This dual approach to the problem of I/O is intended to illustrate the fundamental difference between the design of systems, as opposed to application, routines.

The READ routine would contain the channel programs corresponding to perhaps three or four buffers, together with the buffer management logic to maintain maximum overlap between channel activity and program instruction execution, that is, to keep the buffers full for as much of the time as possible. The buffers do not necessarily correspond to card images, but rather scatter read techniques employing the chain data facility of CCW's could be used to fill field specific buffers. In any case, the interface between the processor subroutine and the READ subroutine (Fig. 4-3) should consist of a user defined macro which when issued in the processor, calls the READ subroutine.

The address of a card image or the addresses of a group of fields from a card image would be returned to the processor in a parameter list. The buffer management routine would, of course, be responsible for synchronizing the EXCP I/O requests with the allocation of buffers.

An alternative approach to the relationship between processor and READ subroutine might be to have the processor specify the address or addresses into which the READ routine is to read the next card image or set of fields. The READ routine would then be required

to create dynamically, through instruction execution, the appropriate channel program. This arrangement implies that the buffer management function is to be included in the processor; otherwise the potential for overlap is greatly restricted.

In either approach the programmer's success in achieving CPU/channel overlap is indicated by the accumulated CPU wait time printed, together with the total execution time, at the end of each SALT job. The relationship between CPU wait time and I/O device characteristics can be explored by successive SALT runs in which the only change is the time parameters in the "DEVICE=" option on the \$SALT card.

The lectures, while assignment #3 is in progress, might be concerned with the following topics:

- (a) the structure of a supervisor as a collection of routines operating in real time in a logical relationship to each other determined by interrupts and by the execution of instructions for branching and loading of new program status words,
- (b) a task (both system and program) represented by a collection of logically connected instructions related to other tasks through the interrupt mechanism,
- (c) resources and their attributes,
- (d) queues and the queuing of requests for a reusable resource,
- (e) task supervision as an example of managing a resource, the CPU.

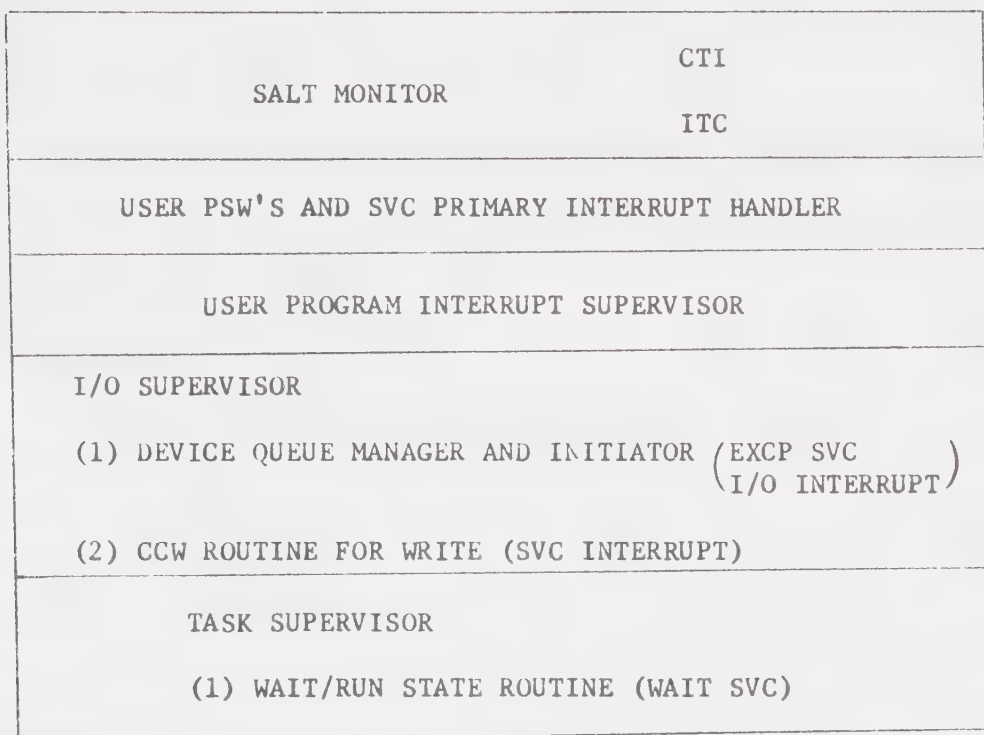
These topics form the background necessary for assignment #4

which consists of transforming the rudimentary interrupt routines and isolated software components of the previous assignment into a unified, embryonic, resident supervisor to service the needs of the problem program of assignment #3.

The structure of a minimal supervisor is shown in Fig. 4-4. Under this arrangement, an SVC primary interrupt handler determines the type of supervisor support requested and calls on the I/O supervisor in the event of an EXCP request or on the TASK supervisor for a WAIT request.

In essence, the I/O supervisor manages a queue of requests for each device (assuming one device on one channel) by enqueueing the CCB of each EXCP specifying the given device, and dequeuing a CCB on the occurrence of the I/O interrupt which indicates the end of the channel program associated with that CCB. Implied in the dequeuing procedure is the initiation of the channel program associated with the next CCB on the queue (unless the queue is empty) by execution of the privileged instructions SIO, TIO and HIO. Housekeeping functions such as the reflection of status information to the appropriate CCB after each I/O interrupt, and handling attempts to read past EOF or EOY belong in the I/O supervisor. The WRITE support routine can conveniently be included with the I/O supervisor although it could as readily be run outside the supervisor state entirely.

For this assignment the TASK supervisor is more rudimentary than the I/O supervisor. Its single function is to place the CPU into the WAIT state with I/O interrupts enabled whenever a WAIT



PROBLEM STATE

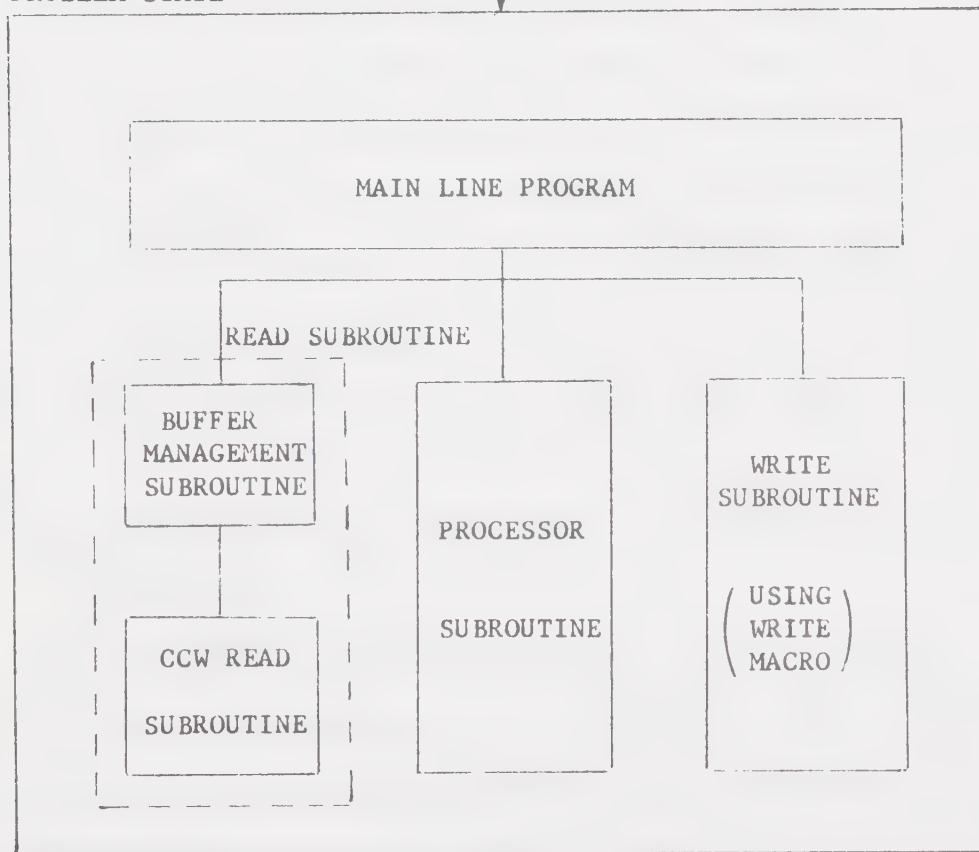


Fig. 4-4. Fourth stage with user task and I/O supervisors

SVC has been executed.

The lectures while assignment #4 is being completed might center on non-supervisor system topics such as:

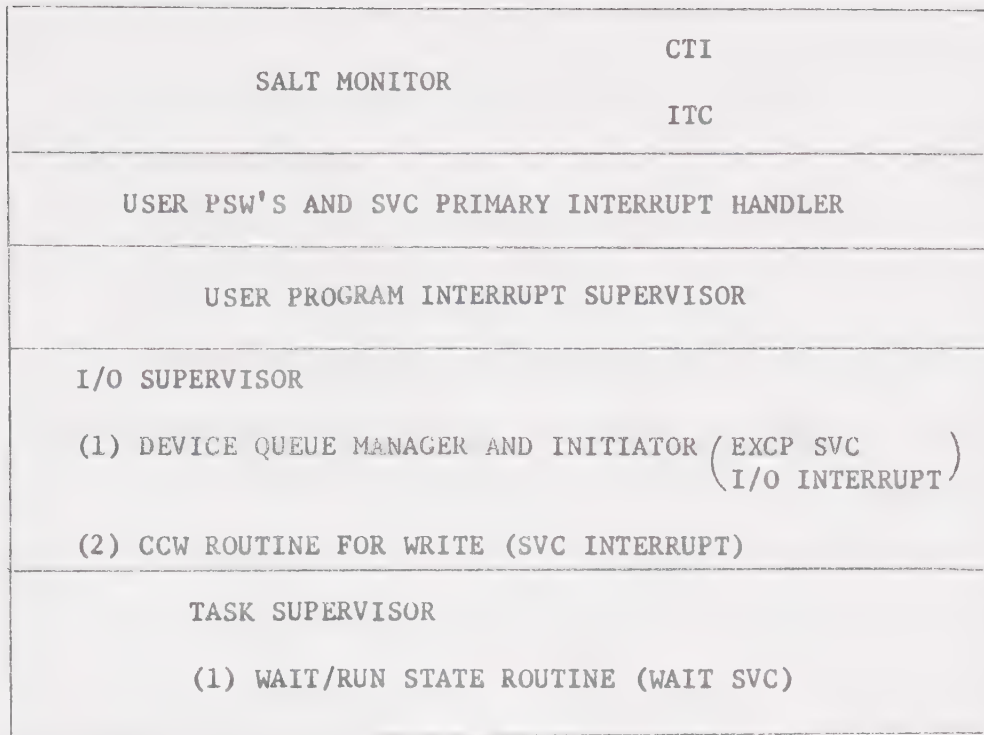
- (a) assembler design,
- (b) compiler construction,
- (c) loader and link-editor functions,
- (d) utilities .

Assignment #5 (Fig. 4-5) is intended to illustrate this area. The problem consists of defining a very restricted assembler language, a subset of the SALT language, and implementing the corresponding assembler as the processor subroutine of assignment #4. The language need not contain the facility for symbolic addressing, thus simplifying the assembly procedure to the following:

- (a) scanning source statements obtained from the READ buffer manager to delimit the mnemonic and operand fields,
- (b) checking fields for possible error conditions and issuing appropriate messages,
- (c) converting mnemonic and operand fields to their SALT machine instruction equivalents,
- (d) producing a simple line-at-a-time listing using the WRITE subroutine,
- (e) loading each machine instruction as it is translated.

The lecture material, while the assembler is being developed, might deal with more advanced software characteristics such as:

- (a) parallel, or asynchronous, processing of a number of tasks,
- (b) spooling of input and output data,



PROGRAM STATE

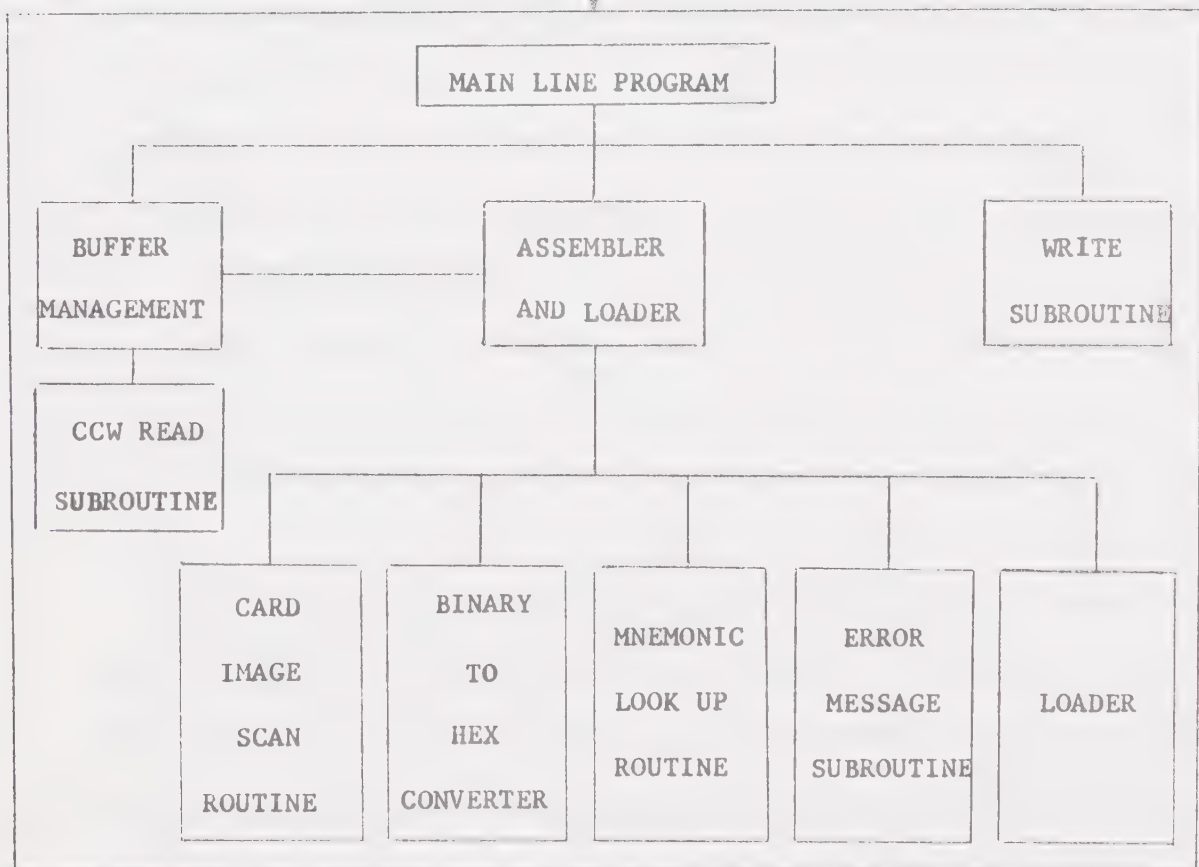


Fig. 4-5. Fifth stage employing user assembler for a greatly simplified assembler language

- (c) spooling of jobs,
- (d) time slicing,
- (e) management of data structures,
- (f) scheduling and management of jobs.

The final assignment in the program (#6, Fig. 4-6) entails a considerable degree of development from #5. A major change to the resident supervisor is required in order to expand the TASK supervisor from its rudimentary form in assignment #4 to include the capability of queuing Task Control Blocks of separately defined tasks to be run asynchronously, and to dispatch these tasks for CPU attention in an attempt to minimize CPU wait time.

An even more basic change is required to re-structure the problem program into a parent task that defines to the supervisor three subtasks:

- (a) an asynchronous reader that spools card images into a large memory work area,
- (b) an assembler/loader that uses the services of a buffer management routine to interface to the other tasks,
- (c) an asynchronous writer that writes to the output device, lines placed in a large memory work area by the buffer management routine at the request of the assembler/loader.

The definition of the tasks to the supervisor, the relationship between tasks (Parent or Offspring) and the synchronization constraints between competing tasks should be accomplished by a set of user defined macros similar to the ATTACH, POST, WAIT, EXTRACT, DETACH group of system macros available under the IBM OS-MVT system.

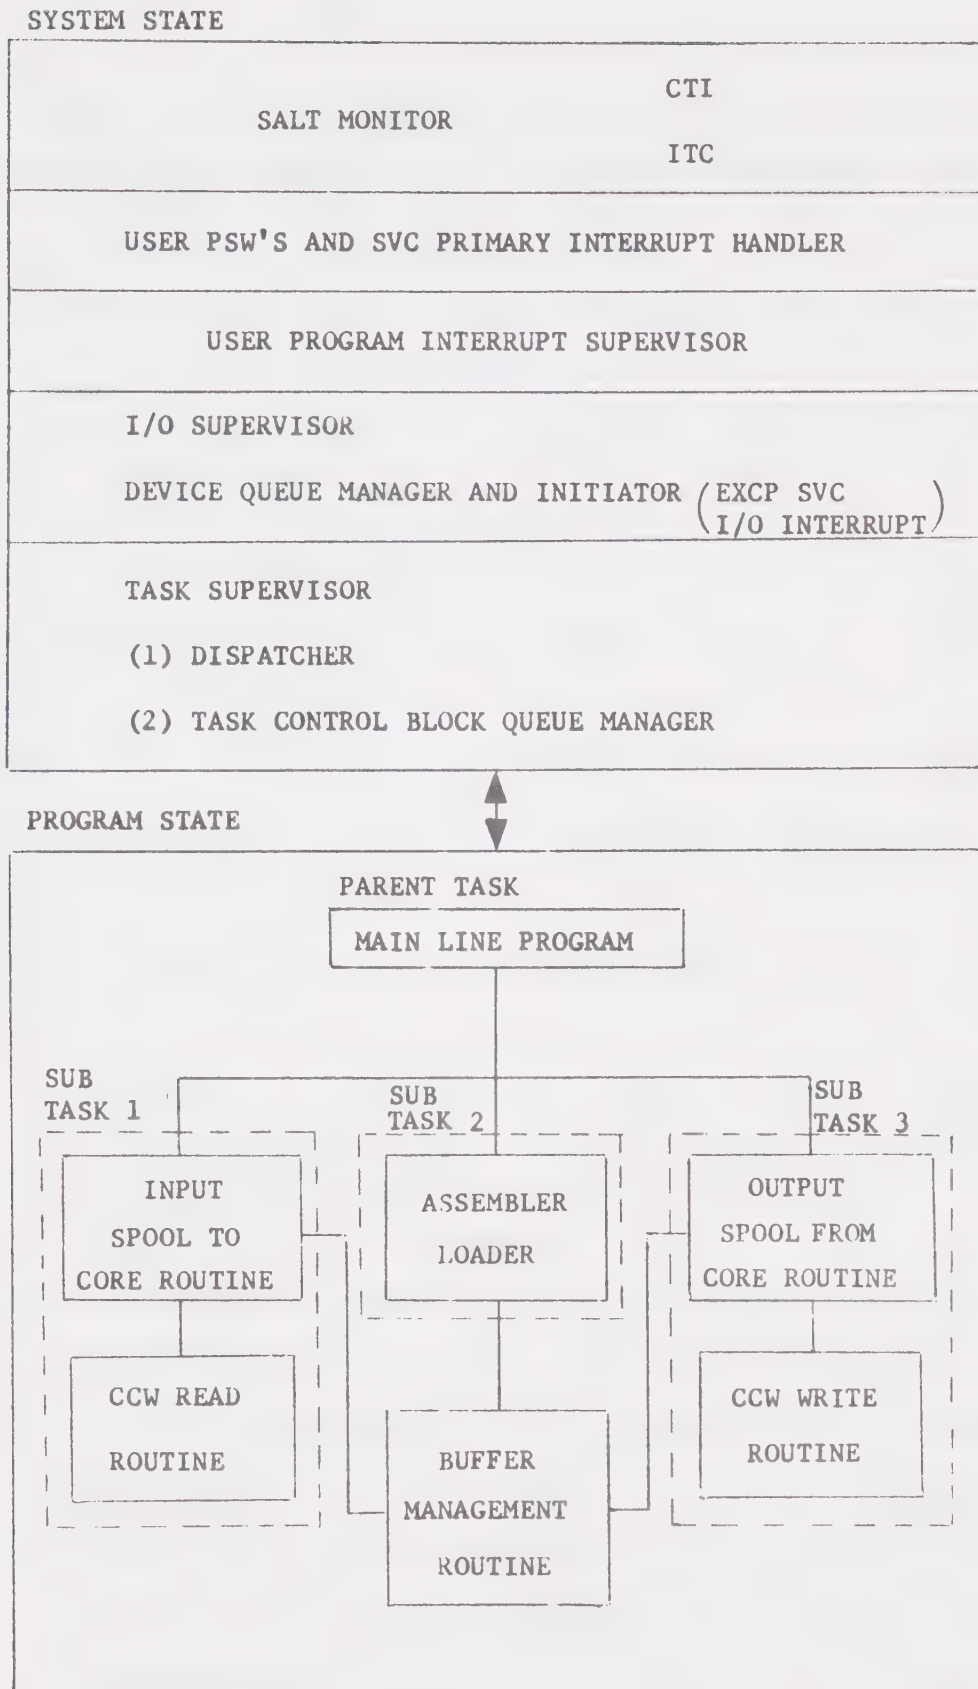


Fig. 4-6. Final stage containing user multiple task supervisor running device/memory spool packages asynchronously with assembler

After the combined system of assignment #6 becomes operational, a number of runs should be made to illustrate the dependency of system performance on the following factors:

- (a) the size of the reader and writer memory work areas,
- (b) the "DEVICE=" option parameters,
- (c) the relative priorities between the reader, writer and assembler tasks.

The six preceding assignments represent a considerable amount of programming, probably too much to expect a student to complete in a one-term course. One solution would be to form groups in which the individuals develop different sections of the system. Another solution that appears to offer more advantages is to pre-program all six assignments and have each student create his own system, using selected routines from the pre-programmed version. The advantages of this are:

- (a) The student is relieved of the work of writing and debugging parts of the system that are not interesting or particularly illustrative,
- (b) Each individual can examine another programmer's work and thus become exposed to new techniques,
- (c) The student is forced to abide by a standardized organization and module interface in order to use the pre-programmed routines, thus ensuring that he eventually develops a system that can be recognized as fulfilling the requirements of assignment #6.

4.4 Summary

The set of assignments outlined in the previous section is not intended to illustrate all possible uses of SUPERSALT, but the author feels they do demonstrate that the major problems associated with using the manufacturer's software can be overcome by using the SALT extension. The assignment descriptions and course outline are intended to describe an entirely feasible systems course plan and could be used as a guide in organizing future systems courses. As such, they reflect what the author believes to be important topics. For example, the hardware is certainly visible to the programmer of the six assignments, and it exists in a form sufficiently simple to enable the user to supply his own supervisor to accommodate his problem program. Hopefully, the user could create an increasingly sophisticated supervisor in response to the increasingly sophisticated requirements of his evolving problem program, thus gaining insight into both the application side and that of the supporting supervisor. At the same time, the lectures would no longer be a simplification, and therefore a flawed representation of a complex, pre-existing system, but rather would use the much simpler system developed in the assignments as a base from which to launch into a more complex and theoretical treatment of software organization.

CHAPTER V

IMPLEMENTATION OF SUPERSALT: THE MAJOR PROBLEMS

5.1 The Internal Facilities Required for CCW I/O Support

In order to properly simulate I/O operations to the user, the SALT system maintains tables containing descriptive and current status information for devices and channels defined, through the \$SALT statement, to be part of the simulated machine. These tables are used to create and to maintain a stack of events, related to I/O activity, that provide the interface between the real input/output in real time and SALT input/output in SALT time.

For each possible device, there is defined a device description table whose format appears in Fig. 5-1. The first, third and fourth fields are inserted into the table from the fields of the 'DEVICE=' option (Appendix II) on the \$SALT statement when it is scanned by the job control processor (JOBCTRL) at the start of a SALT job. The second field contains a device type code used to determine the physical characteristics corresponding to this device. The device description tables are used by the system during execution of a SALT program to determine:

- (a) if a unit on which I/O is requested actually exists,
- (b) when, in SALT time, an initiated I/O operation should begin transmitting data across a channel,
- (c) when, in SALT time, an initiated I/O operation should terminate.

The channel description table consists of a full word for each

of the eight possible channel designations (0 to 7). Each word contains a number indicating the "width" of the corresponding channel, that is, the number of bytes that can be transmitted across the channel in parallel. This information is used with the record length implied by the device type code in order to calculate the number of "pulse fronts" required by an I/O request that are to be serially transmitted across the channel.

For each device description table there is a device status table whose format is given in Fig. 5-2. The status tables reflect the instantaneous status of each I/O device within SUPERSALT at any given time. With the exception of the pointer to the CCB for this I/O operation, each field in the device status table is brought up to date after each pending event. That is:

- (a) The second field is changed to point to the next (chronological) pending event for this I/O.
- (b) Any change of status is reflected to the status field.
- (c) The residual count field is decremented if information has been exchanged between a SALT program and an internal system buffer.
- (d) The CCW address field may be changed if the current event specifies the end of a CCW which is chained to the next.

The channel status indicators consist of a double word with one byte reserved for each channel. A non-zero value implies that the corresponding channel is currently involved in an I/O operation.

The description and status tables are used in interpreting user-written CCW's. The problem can be adequately illustrated by

CHANNEL DEVICE ADDRESS	$\left\{ \begin{array}{l} 00\ 00 \\ 04\ 00 \end{array} \right\}$
DEVICE TRANSMISSION RATE	
AVERAGE DELAY TIME BEFORE TRANSMISSION	

Fig. 5-1. The device description table

CCB ADDRESS	
NEXT PENDING EVENT ADDRESS	
STATUS FLAGS	RESIDUAL COUNT
CCW ADDRESS	

Fig. 5-2. The device status table

considering a user CCW employed to perform a simple card read. The read is initiated when the SVC associated with the EXCP macro is executed, causing a transfer of control into the SALT monitor. (If the user is handling his own interrupts, control passes to his I/O initiation routine and the I/O does not actually begin until the SIO instruction within that routine is executed.) After checking the validity of the command, device address, channel address and receiving area address, the monitor initiates an actual

read operation on a /360 channel. No further /360 instructions are executed within the SALT system until the read is finished.

The completion of the actual I/O is hidden from the SALT program, and SALT instruction execution is resumed. At a time synchronized to SALT instruction execution according to the device characteristics, "channel width" portions of information are "transmitted" from the internal system buffer to an address calculated from that specified in the CCW. It is the "remembering" of these "transmission" events that forms the most interesting aspect of the implementation. Each event consists of fields specifying:

- (a) the time at which the event will occur,
- (b) the type of event that will occur,
- (c) the addresses between which data will be transferred,
- (d) other control information,

and forms an element in a data structure. The type of data structure chosen was determined by the following requirements:

- (a) The elements in the structure must be logically ordered in chronological sequence.
- (b) Since previously initiated I/O may be in progress when a given I/O operation is initiated, the data structure must allow the interleaving of new elements "between" old elements to preserve the ordering.
- (c) As events "happen" they should be easily removed from the structure and the memory space they occupied made available for re-use.
- (d) Since channel execution of a CCW can be halted before nor-

mal termination (HIO), it must be possible to remove efficiently from the structure all events associated with a given channel program even though they may not be logically adjacent, that is, events for other I/O may intervene.

The above requirements clearly dictate that the data structure be a form of linked list with a main linkage that links the events in chronological sequence and a secondary linkage that links events for the same CCW from first to last. Since events may be of several types requiring different amounts of information, the list elements must be variable in size.

The memory space for the data structure is dynamically allocated from the high end of the region at the beginning of execution of each SALT job, and thus reduces the space available for the machine instructions of SALT programs. It is therefore necessary to make the most efficient use of the space allocated to the data structure. A traditional method [12] of managing space allocation for a linked list is, in effect, to keep two lists; the second linking the free space interspersed within the allocated elements. As new elements enter, the list space for them is obtained from the free memory list and as elements are deleted, the space they occupied may be immediately returned to the free memory list or left for collection by the later execution of a garbage collector.

Because of the extremely active nature of the data structure when a SALT program with large I/O requirements is run, the garbage collection method was eliminated from consideration. It was felt

the collection routine would be called to rebuild the free storage list too frequently in the restricted memory space, resulting in high overhead. The other method of attaching elements to the free storage list as they were deleted would ordinarily be suitable, but some problems are encountered in the case of a list with variable length records in that list elements of differing sizes are required from, and released to, the free list. In order to keep memory fragmentation to a minimum, the free list would need to be maintained in order of size of element from smallest to largest, and some scheme should be implemented to consolidate two or more adjacent elements into one free list element. For a highly active structure, this procedure entails a considerable amount of searching lists and relinking of elements into and out of them.

An alternative method to these traditional solutions was employed to minimize overhead. This method involves restricting the variable sized elements to multiples of eight bytes and employing an allocation control block in which every bit indicates by its value the current status of a particular eight byte field in the data structure (the pending event stack). That is, if the third bit is one, then the third double word in the data structure is currently being used. The procedure for allocating space for an element requiring, for example, three double words is as follows:

- (a) Starting at the left of the allocation control block, scan toward the right, looking for the first three consecutive bits of value zero.
- (b) When these are found (assuming they are), set them to value

one and calculate the corresponding address in the pending event stack as follows:

$$\text{address of free element} = (\text{number of bits offset into} \\ \text{the allocation control block} \times 8) + \text{address} \\ \text{of pending event stack}$$

- (c) In the newly allocated element save the one byte mask field used to "OR" the allocation bits to one to facilitate the eventual freeing of the space; the element can then be created and logically linked into the stack.

The calculation of the free element address can be performed very efficiently in a /360 since the value of the offset ends up in a register. The multiplication by eight is a left shift of three positions, and the addition of the result with the stack address can be achieved with a simple "load address" instruction. However, designing a procedure for efficient scanning of the allocation control block for consecutive zero bits presented a considerable challenge.

The procedure for de-allocating the space occupied by a newly deleted element consists of:

- (a) calculating the number of bytes offset into the allocation control block required in order to use the one byte mask stored during allocation.

$$\text{number of bytes offset} = (\text{address of this event} - \text{address} \\ \text{of start of stack})/64$$

Division of 64 can be achieved by a right shift of six positions.

- (b) "exclusive ORing" the mask byte with the appropriate allocation control block byte to set the bits corresponding to the memory reserved for the element to zero, freeing the area for re-allocation.

The layout of a typical pending event is shown in Fig. 5-3. This event controls transfer of information between the system buffer and user area. Other event types have a similar format and all have identical fields in the first sixteen bytes. Other event types and their uses are:

- (a) Control - to cause, at the device, some form of mechanical activity other than reading or writing records, for example, spacing the printer,
- (b) I/O Interrupt - to cause a SUPERSALT I/O interrupt to mark the end of a channel program,
- (c) Chain Command Marker - to mark the last event of a CCW that had the chain command flag on, and to contain the address of the CCW "chained" to,
- (d) Chain Data Marker - to mark the last event of a CCW that contained a one bit in the chain data flag and to point at both the next CCW to be processed and at the position in the device buffer to which, or from which, the I/O operation is to continue.

0	2	3	4	8	9	12	13	14	16	20
LINK	NO. OF BYTES	TYPE	TIME	NXT EVENT	CCB ADDR	COM	ALLOC MASK	STATUS TAB DISP	FROM ADDR	TO ADDR

LINK: contains the displacement to the next chronological event

NUMBER OF BYTES: specifies the number of bytes of information to be transmitted by this event

TYPE: indicates by its value that this event is to cause a transfer of information

TIME: specifies the SALT time at which this event is to occur

NXT EVENT: contains the double word displacement to the next event in the sequence for this CCW

CCB ADDR: points to the CCB associated with this event

COM: contains the command code for this CCW

ALLOC MASK: contains the bit pattern required to free the memory occupied by this event after the event has "occurred"

STATUS TAB DISP: contains the displacement to the status table for this event

FROM ADDR: specifies the address from which "NO. OF BYTES" of information is to be transmitted

TO ADDR: specifies the address to which "NO. OF BYTES" is to be transmitted

Fig. 5-3. The fields of a pending event to control information transfer

The correspondence between a typical CCW and the events that represent it are shown in Fig. 5-4. As a result of the execution of the SVC for the EXCP, the SALT system calls on the CCW processor which inspects the CCW and determines that:

- (a) Two information transfer events of twenty bytes each would exhaust the count field of the CCW.
- (b) The first event should be placed in SALT time "Average Delay Time" microseconds beyond the current SALT time.
- (c) The other events should follow at $(\text{channel width} \times \text{constant}) / (\text{data transfer rate})$ intervals.
- (d) The last event should indicate that chain data was requested, and point at the CCW chained to.

Following creation and insertion of these events into the Pending Event Stack, the SALT system continues with SALT instruction execution. After execution of each instruction, an appropriate value is added to the accumulated SALT time. Between instructions, the Pending Event Stack is checked to see if the next event is scheduled to "happen" now. In the example, the first event would be "accomplished" by moving twenty bytes from the system buffer to LINE, while the third event would involve moving twenty bytes to LINE + twenty from the system buffer address + twenty. The fourth event would involve the fetching of a new CCW and the calling of the channel processor to process it. The new CCW contains the address into which the READ is to continue while the current position in the system buffer would be recorded in event four.

From the above example it is seen that only one CCW from a chan-

CCW1 CCW READ,CARD,X'80',40

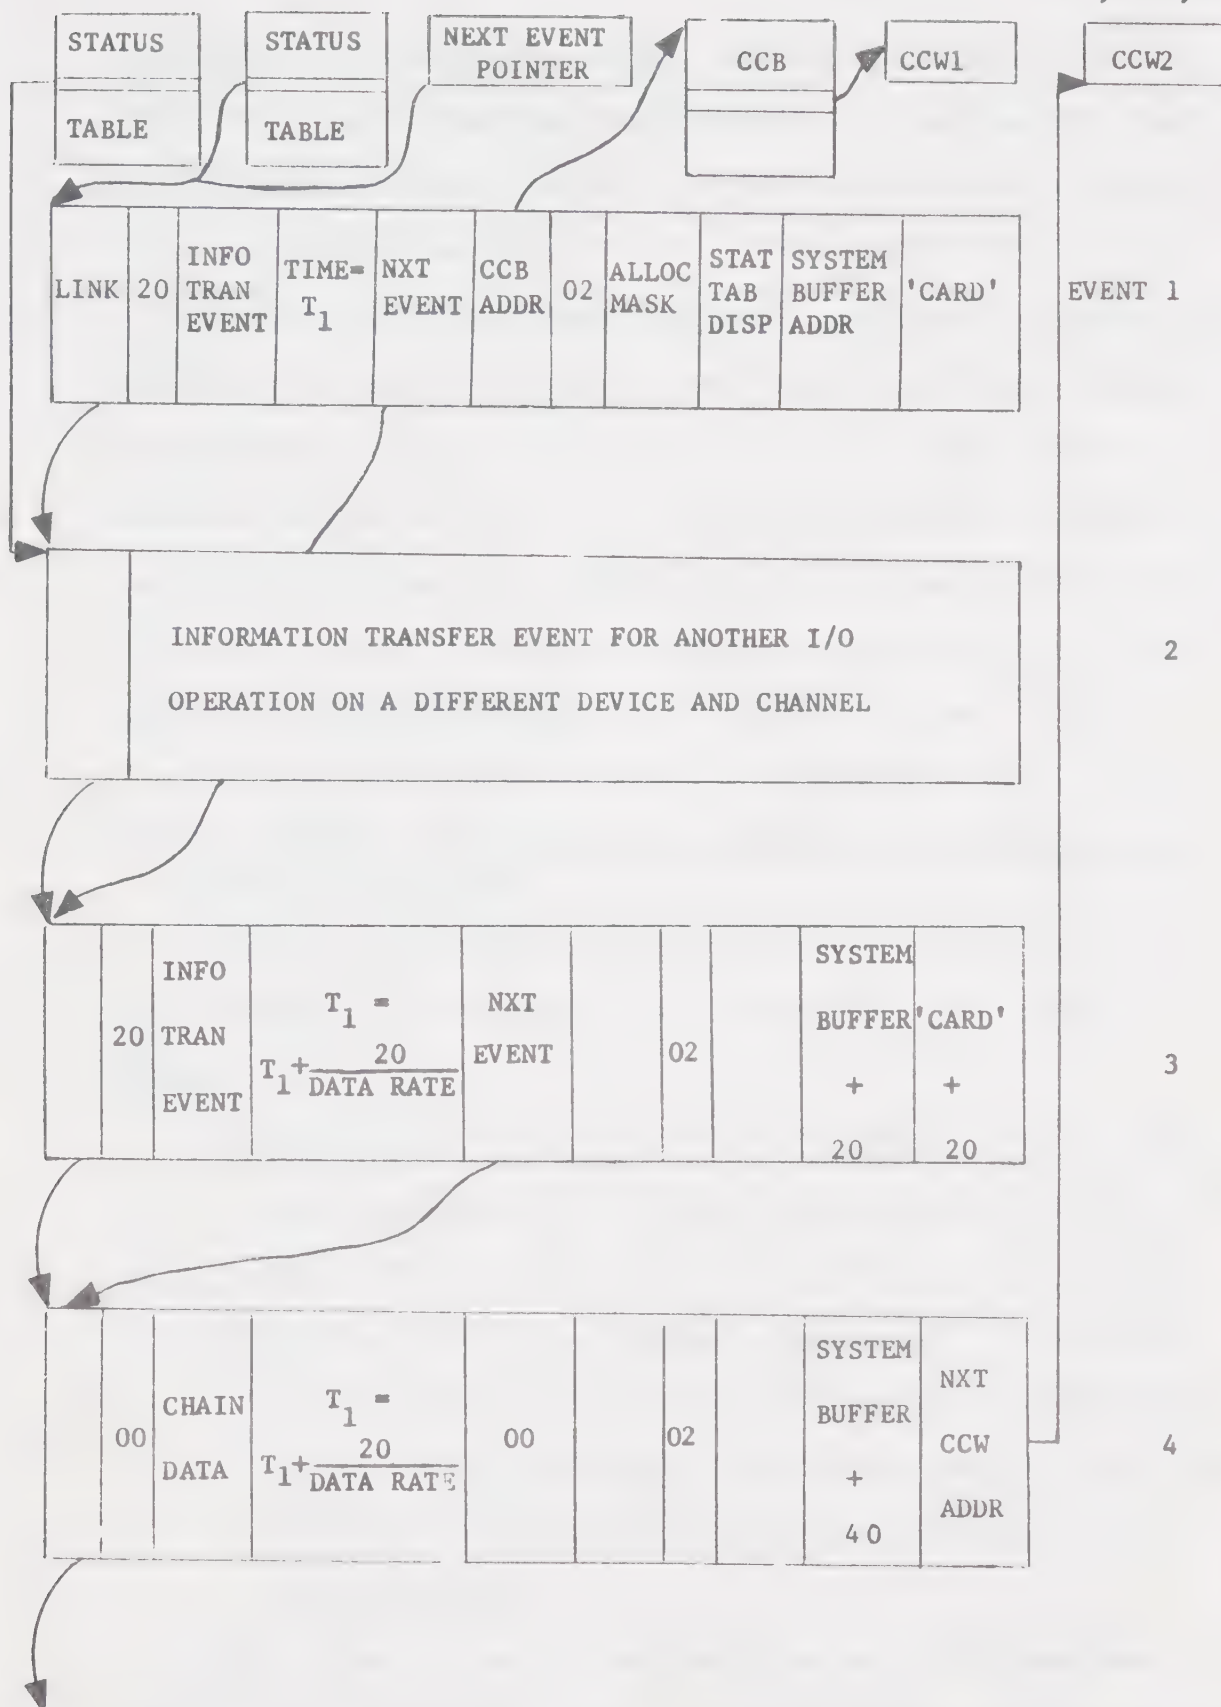


Fig. 5-4. Pending events corresponding to a CCW initiated at time T_1 minus "Average Delay Time"

nel program is processed at a time. This is necessary, both to reduce the number of pending events maintained at any one time and because the execution of one CCW in a chain could be used to read in the next CCW in the chain.

In practice, the preceding scheme has worked well for a machine configuration including one eighty byte record input device and one 120 byte record output device. With twenty byte channel widths, the worst case situation is ten pending data transfer events plus two pending I/O interrupt or chain events which require a total of 272 bytes. In fact, a double word is reserved for allocation bit indicators, implying a present stack size of 512 bytes.

5.2 The Internal Facilities Required to Simulate a Multi-State CPU with Interrupt System

Since SALT instructions are interpretively executed rather than executed on /360 hardware, the simulation of a CPU can be achieved by:

- (a) monitoring the state of the CPU as presented in the current status indicators to determine whether instruction execution should proceed, and, if so, what instruction set is valid,
- (b) creating a SUPERSALT interrupt whenever an SVC instruction is executed, or between instructions whenever an I/O interrupt type of pending event is serviced or whenever an invalid condition is detected during interpretive execution of a SALT instruction.

If the "PSW=" option has been specified on a \$SALT statement, then the SALT interpreter, prior to SALT instruction execution, moves

the first eight bytes from the address given in the operand field of the option to an internal field used as the current program status word (PSW). This process is somewhat similar to the /360 initial program load procedure. If the initial PSW specifies "run" state, then SALT instruction execution proceeds from the address specified in the PSW until the first interrupt. An interrupt is simulated by:

- (a) moving the current PSW field to the Old PSW slot,
- (b) moving the New PSW slot to the current PSW field,
- (c) moving the SUPERSALT general purpose register fields to the Old GPRS slot,
- (d) moving the New GPRS slot to the SUPERSALT register fields,
- (e) proceeding with instruction execution (if run state has been entered) under control of the new current PSW.

The WAIT state is simulated by ceasing instruction interpretation until the next SUPERSALT interrupt can be forced to occur by advancing the SALT time scale. This is done through clearing any pending events from the stack, one after another, advancing SALT time appropriately in each instance until an I/O interrupt on an unmasked channel is encountered. If none is found before the stack is exhausted, the program is terminated for running overtime in a locked WAIT condition.

The process of masking a channel and thus postponing any I/O interrupts from that channel is simulated by leaving the appropriate I/O interrupt event in the stack unserved past its scheduled time. Between each instruction an attempt will be made to

service this interrupt. The attempt will be successful if the mask has just been removed; otherwise the next event in the stack is checked to see if it is scheduled to occur now. A masked-off I/O interrupt can therefore remain at the head of the stack while events that become current are serviced from "behind" it.

Implementation of routines to interpret privileged SALT instructions presented few difficulties and they are only briefly described here.

The routine for SIO simply calls the CCW processor and provides a dummy internal CCB. It uses the return information from the CCW processor to set the condition code and status table entries, and may move status information from the device status table to the user-supplied Channel Status Word.

The routine for HIO extracts the address of the first pending event (if any) for the specified device, releases the memory for it and re-links its neighbours, uses the NXT EVENT field to calculate the next pending event in the stack for this device, deletes it and continues following the chain until the last pending event for this device has been removed from the stack. The routine then resets the status table, thus releasing the device for subsequent I/O operation.

The routine for TIO instruction extracts status information from the device status table and inserts it into the channel status word.

The routine for the SSM instruction simply resets the system mask field of the field containing the current SUPERSALT PSW with a

specified bit pattern.

The routine for the LPSW instruction moves the sixteen words from the New GPRS field provided by the user to the SUPERSALT register fields in the interpreter and moves eight bytes from the address specified to the current PSW field.

5.3 Variations in the System Format

The extensions to the system were implemented in a form that uses the conditional assembly feature of the /360 Assembler Language. By appropriate settings for two symbolic parameters and subsequent re-assembly, the following versions of the SALT system can be generated from the same source program:

- (a) The fully extended system: this version has support for channel programming and user interrupt handling. If both the "DEVICE=" and "PSW=" options are employed in a job, then that job is restricted to a maximum about 10K smaller in object size than under the original SALT system, and that job may require up to twice (depending on I/O demands) the /360 CPU time for SALT program interpretation. If only the "DEVICE=" option is specified, then the CPU overhead is slightly reduced. If only the "PSW=" option is specified, implying that only READ's or WRITE's are used in the program, then CPU overhead is only slightly degraded from the original SALT system.
- (b) The semi-extended system: this version has support for channel programming. The "PSW=" option is not recognized

as valid on the \$SALT card and no internal current PSW is maintained in the system nor is an interrupt mechanism simulated. The CPU overhead is slightly less than using (a) with the "DEVICE=" option only, while memory size restrictions are only slightly improved over (a).

- (c) The basic system: this system has no support for user interrupt handling or channel programming. The "DEVICE=" and "PSW=" options are both invalid on the \$SALT card. The user sees the system as described in Easton and Penny [10], and internally it is little changed from the system described in Dutton [7]. The CPU overhead and memory requirements are exactly the same as those of the original SALT system.

In accordance with good software design principles, the three systems are compatible in one direction. That is, a program designed for and run on system (c) can be run unchanged on systems (a) and (b). Similarly, a program run on system (b) can be run unchanged on system (a).

CHAPTER VI

LIMITATIONS OF SUPERSALT AND PROPOSALS FOR ITS
FUTURE DEVELOPMENT6.1 Limitations

A software system can be evaluated by examining the extent to which it fulfills its intended roles. The original aim in creating SUPERSALT was to produce an operational, distributable system such that:

- (a) The user would be convinced that he is programming a complete computer, that is, one in which there are no inconsistencies in the time base.
- (b) The user would be aware of the characteristics of, and responsible for the activity of, peripheral devices and channels operating asynchronously with the CPU.
- (c) The user would have the facility for developing within the simulated machine, his own software systems on virtually any level of complexity from basic interrupt handlers for selected types of interrupts to full-scale operating systems.
- (d) The user would be able to explore the relationships between various machine configurations and the software characteristics necessary to utilize the hardware effectively.
- (e) The above characteristics would be supported in a system that protects the installation from the user and makes modest demands for memory space, /360 CPU attention and /360 I/O resources.

The following discussion points out that these aims have not been entirely fulfilled, and Section 6.2 describes extensions to SUPERSALT designed to overcome these shortcomings. The proposals in the next section were not implemented in the original extension in order to allow time to evaluate, through experience gained in using SUPERSALT, the precise form that future developments should take. However, for each development proposed, a general plan of implementation is provided to illustrate its feasibility.

The simulated computer system, specified in Chapter III, appears to meet the requirements stated in (a) and (b). Close examination reveals, however, an inconsistency in time that is apparent only when one considers the complete SALT job consisting of assembly and execution. Essentially, the problem is that the SALT source program is processed by an assembler running on a /360 machine while the resulting SALT object program is executed immediately afterward on a simulated, simplified /360 (SUPERSALT) which executes instructions in the order of 100 times slower than the actual /360 Model 65. In other words, a SALT job is completely assembled and loaded into core, and only then is the initial PSW inserted into SUPERSALT to cause its activation. The single SALT object program might contain a supervisor for the SUPERSALT CPU and several programs recognized by that supervisor. The termination of the run is caused by execution of an SVC 0 while in the supervisor state, which appears to deactivate the SALT computer and, in fact, causes the SALT system to perform job-to-job transition.

One possible way of viewing this situation such that the incon-

sistencies become more acceptable is that there are really two computers, a SUPERSALT CPU and a /360 CPU, both sharing the same memory. The assembler is assumed to run on the /360, assembling a SALT job. After loading it into memory, the /360 gives a "tap on the shoulder" to the SUPERSALT CPU, causing the initial PSW to be loaded and the SALT program executed. The execution of an SVC 0 on the SUPERSALT CPU while in the supervisor state is assumed to cause a "shoulder tap" back to the /360 which has been in the WAIT state, and then assembly of the next job proceeds. Similarly, other supervisor functions reflected to the SALT monitor can be considered as receiving service from the /360 CPU.

The aims expressed in (d) and (e) above are only partially realized in SUPERSALT. It is certainly true that the user can create his own interrupt handler to explore the operation of the multi-state CPU as is done in Appendix V. However, it is certainly not true of the extension described here that the user can create his own full operating system to virtually any level of complexity. Those functions of an operating system that can be demonstrated by taking an instantaneous view of the operation of the system can, in general, be accommodated by SUPERSALT. A typical example of this type of function is the creation of asynchronous tasks within a job. It is quite feasible to create a supervisor that supports, and a problem program that requests, subtasking, and have the supervisor and program assembled and executed as the same SALT job. In general, only those functions that can be illustrated as operating in an interrupted fashion over a comparatively long time scale are not

adequately supported by SUPERSALT. Examples of this type are scheduling and initiating jobs and passing information from job to job through data sets.

For example, it is quite possible as outlined in Chapter IV, to create a software system that contains a supervisor which runs an assembler and an input stream of pseudo programs for assembly and execution, all within one SALT job. This arrangement implies that modularity can only be achieved through separate tasks connected by PSW exchanges or through subroutines of the same assembly. It is not feasible to:

- (a) cause the output of one execution of the assembler to be combined with a subsequent output to form a single object program,
- (b) create a data set in one program that is to be used by a subsequent program,
- (c) maintain libraries of data sets global to all programs within a SALT batch,
- (d) maintain libraries of data sets global to all SALT batches,
- (e) cause the input to or output from a SALT job to be spooled between devices,
- (f) have student designed data set organization schemes used within SALT jobs.

Clearly, if some form of pseudo device were supported which allowed read/write data sets whose records were independently or nearly independently accessible, the limitations mentioned above could be largely overcome. Further, if this device were interfaced to /360

direct access devices, the shortcoming mentioned in (d) above could be eliminated.

The final aim, (e), mentioned at the beginning of the chapter has been met. The installation is fully protected from the user, SUPERSALT will run in 100K of memory or even less for small SALT programs, /360 CPU requirements (even though perhaps doubled from the original system for some jobs) are still many times less than would be the case using the manufacturer's software for the same purpose, and /360 I/O requirements have not increased over the original SALT system. For example, the sample program in Appendix V was run in a region of 100K and was assessed a total charge of \$.25 by the charging system in use at the University of Alberta under the OS/MVT system.

6.2 Recommendations for Further Development

It is apparent from the discussion in the previous section that the largest potential improvement in the system could be obtained by providing some form of simulated direct access device that is programmable by the user. The internal organization of the system as set out in Chapter IV does not preclude the addition of such a device, providing the characteristics of the I/O are carefully defined with a view to simplicity. One such definition follows.

Records, the unit of information, are individually addressable on such a device through specifying a compound address, the parts of which can be considered as specifying the nodes of a tree structure representing the physical organization of the device (Fig. 6-1).

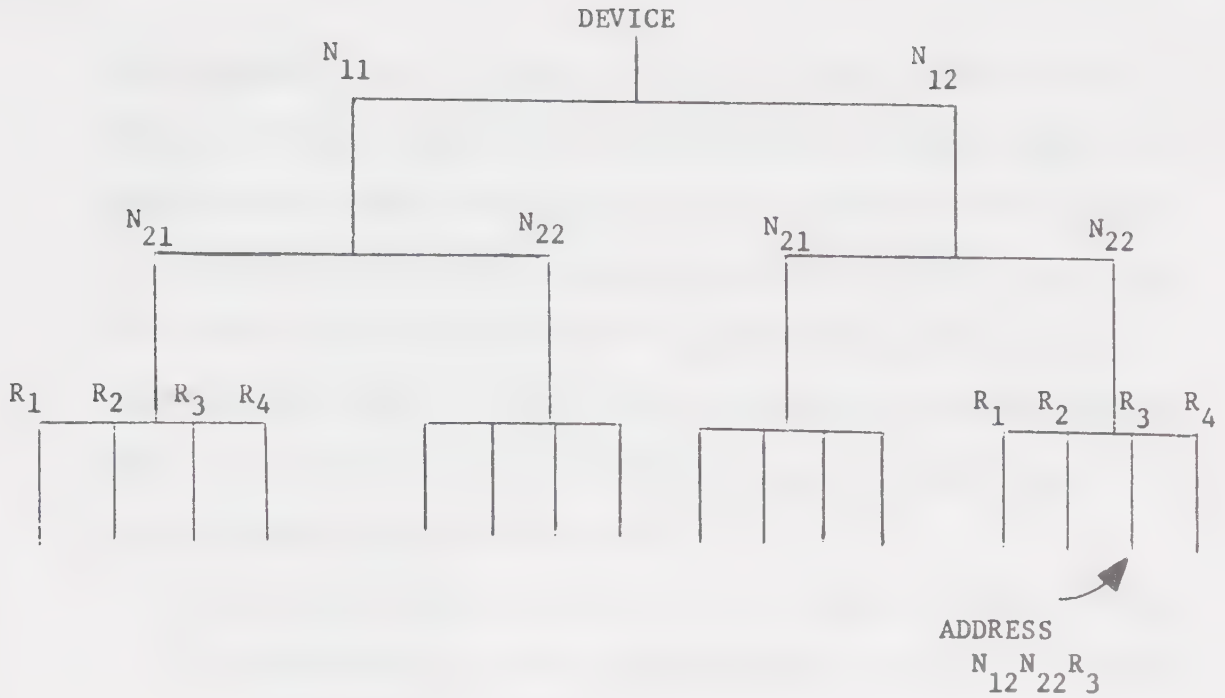


Fig. 6-1. A direct access device addressing scheme

The device is assumed to have physical mechanisms that require the channel, through execution of CCW's, to cause explicitly the shift from one node to another, either down or across. For example, if the last record read, using the notation of Fig. 6-1, was $N_{11}N_{21}R_4$, the record $N_{12}N_{22}R_3$ could be read by the following channel program using command chaining:

- (a) a CCW to place the mechanism at node N_{12} ,
- (b) a CCW to place the mechanism at node N_{22} ,
- (c) a CCW to place the mechanism at record R_3 ,
- (d) a CCW to cause the record "pointed at" to be read.

If a second READ were issued without repositioning the mechanism, the next record, R_4 , would be obtained. A third READ would retrieve record $N_{11}N_{21}R_1$. If a READ or WRITE were issued in which the count

field exceeded the record size, only that record would be read or written, and the status would indicate that the record length was exhausted before the count field, with the number of bytes remaining in the count specified by the residual count field. In brief, this device is a unit record device in which the records are re-readable, re-writable and individually addressable. Ideally, the device is assigned through the "DEVICE=" option, describing not only the device address and type but:

- (a) the number of bytes in a record and the data transfer rate,
- (b) the number of records grouped under a N_{2-} node and the time required to position at a record after positioning at the N_{2-} node,
- (c) the number of N_{2-} nodes grouped under a N_{1-} node and the time required to position at an N_{2-} node after positioning at the appropriate N_{1-} node,
- (d) the number of N_{1-} nodes and the time to proceed from one to another.

The implementation of such a device seems to be entirely feasible using the current internal organization specified in Chapter V. The device characteristics specified on the \$SALT card could be kept in an extended device description table and used by the CCW processor in managing the Pending Event Stack. The status table would require an extension to keep information on the current position of the read/write mechanism. The pseudo device could be interfaced to the real installation direct access devices through the SALT system maintaining an actual data set whose records each contain,

say, a complete record group under a N_{2-} node and are individually accessible. The procedure for simulating the set of CCW's outlined above for reading record $N_{12}N_{23}R_3$ after reading $N_{11}N_{21}R_4$ would be:

- (a) Process the first CCW and enter two pending events into the stack, one to cause a change of address in the status table from $N_{11}N_{21}R_4$ to $N_{12}N_{21}R_1$ and one to cause command chaining, and flag the status table indicating a disk read will be required.
- (b) When the chain command event is serviced, process the second CCW and enter two pending events into the stack, one to change the mechanism from address $N_{12}N_{21}R_1$ to $N_{12}N_{23}R_1$ and the other to cause command chaining, and ensure the status table is flagged for a disk read.
- (c) When the chain command event is serviced, process the third CCW and enter two events into the stack, one to change the mechanism from address $N_{12}N_{23}R_1$ to $N_{12}N_{23}R_3$ and one to cause command chaining.
- (d) When the chain command event is serviced, process the fourth CCW and since it specifies READ and the READ flag is on in that status table, read into a system buffer the actual disk record corresponding to $N_{12}N_{23}R_3$, calculate the address of R_3 and move it to another system buffer, and insert the necessary events into the stack to move the required number of bytes of the buffer to the SALT program followed by an I/O interrupt event to mark the end of the channel program.

Even though the implementation of a device like the above seems

entirely feasible in the context of the extended system structure, several hazards are apparent from an operational point of view.

They are:

- (a) If a user specifies an unreasonably large record size or an unreasonably large number of records within each N_2 -node, then a large SALT system buffer will have to be reserved, and for the case of large record size an unreasonable number of pending events will have to be accommodated on each I/O operation.
- (b) A user may, through carelessness or error, cause more /360 disk operations to be performed than are necessary, thus tying up system resources,
- (c) A user may, through appropriate /360 JCL, cause disk data sets created during a SALT program to be kept almost indefinitely.

Clearly, internal limits on record size and numbers must be set to minimize overhead, and adequate control measures over disk space and usage maintained to avoid contradicting aim (e) stated at the beginning of the chapter.

A second addition to the system which would perhaps not provide quite as much improvement in the flexibility of the system as the direct access support just discussed, would be a provision to run parts of a SALT program on different SUPERSALT CPU's sharing a common memory and performing in a multi-processor configuration. To the user it would appear that he was programming two or more CPU's, each of whose reserved core area was specified on the SALT statement

in a separate "PSW=" option field, all of which may have access to any of the channels or devices described in the "DEVICE=" options and all of which appear to run asynchronously with, and semi-independently from, each other. Fully independent operation, of course, would defeat the purpose of the addition since the illustrative aspect of developing software for multiprocessing systems is controlling the communication and "sometimes" synchronization between all of the CPU's through their interrupt system in such a way as to use the total resource efficiently. Asynchronous operation would naturally not be the case whenever there was memory contention between CPU's, that is, no CPU could fetch (either for instruction execution or during instruction execution) or store a byte of memory that was simultaneously being fetched or stored by another CPU. The implementation of a "Test and Set" instruction like the /360 instruction of the same name would seem advisable to allow control of a common resource by the multiple CPU's.

The implementation of this second addition again seems entirely feasible and depends upon making the SALT interpreter re-entrant. Separate current PSW's, register fields and instruction address registers would have to be kept for each SUPERSALT CPU and the interpreter would have to be "called" to interpret one instruction from each CPU instruction stream in turn. If the current PSW of a particular CPU specified the WAIT state, then no instruction would be interpreted in that CPU in contrast to the other still functioning CPU's, providing the impression to the external world of asynchronous operation. Only one device description and device status block per

device specified on the \$SALT statement would be maintained, since all CPU's share the same physical devices. One undesirable characteristic does appear if, to maintain realism, all the CPU's are considered to operate in the same SALT time, that is, if there is only one pending event stack containing the events for all CPU's, and these events are made to happen along a time base tied to instruction execution time in all the CPU's. This implies that for "n" CPU's, each having an instruction interpreted in turn, the single SALT time base must be incremented by the execution time of the longest instruction of the "n" executed. The hardware can be considered as having a common clock pulse generator that causes all the CPU's to perform instruction fetching simultaneously. Any CPU that finishes executing an instruction before the others must await the next common instruction fetching cycle.

This unusual characteristic, although not very realistic, does not appreciably detract from the advantages offered by the opportunity to develop software systems to control a multiprocessor configuration. However, there is a strong probability that aim (e) stated at the beginning of the chapter would, with such a system, be only partially obtained. Almost certainly the SALT programs would become so much larger that added demands would be made on /360 memory space, and the additional overhead necessitated by a re-entrant interpreter coupled with the huge number of instructions that might be interpreted for multiple CPU's would cause a dramatic increase in the amount of /360 CPU attention required.

A third addition to the system might involve the addition of sup-

port in the SALT monitor for asynchronous subtasks within a SALT job. This would require the inclusion of system macros with the following functions:

- (a) Initiate Task (ITASK) - indicates to the SALT monitor that the address of the task control block specified in the operand contains information necessary to define an asynchronous subtask as an offspring of the originating parent task,
- (b) Task Control Block (TCB) - a table of information relevant to a subtask. The information includes the entry point address, the latest register contents whenever the task is waiting, and fields used by the WAIT and POST macros,
- (c) Wait (WAIT) - allows a subtask or originating task to wait on the posting of a completion of an event from the subject task as indicated in its TCB,
- (d) Post Completion of an Event (POST) - allows a subtask to communicate the completion of an event to all tasks waiting on it,
- (e) Post Completion of a Subtask (RETRN) - provides a means of removing a subtask from active competition for CPU attention.

The above three additions compose what the author feels would be the most valuable extensions from the standpoint of instruction potential. It should be stressed, however, that the first two would have to be implemented carefully and used prudently to avoid wasting the /360 resources.

CHAPTER VII

CONCLUSION

While the study of systems is beginning to be recognized as an important part of Computer Science curricula [2], little attention has been given to the problems of illustrating lecture material through assignments. As was described in Chapter II, the typical student-oriented low level language processor lacks the realistic time base and simulated CPU characteristics to act as the vehicle through which assignments could be undertaken. The author's experience in using the manufacturer's software as the basis on which to set assignments merely demonstrates how efficiently the manufacturer's hardware/software combination isolates the user, preventing him from directly viewing the very aspects with which lectures on systems are concerned. The most direct solution, that of presenting each student with real hardware for which he must develop software, suffers from (in addition to economic problems) a surfeit of unenlightening "real life" complexities and details which would make assignments even more difficult and time-consuming without adding to their worth as a teaching aid.

The alternate solution, that of extending a student processor in order to present to the student the appearance of a simplified computer system, has been investigated from the point of view of:

- (a) the essential characteristics of such a system for purposes of illustration,
- (b) its usefulness as a mechanism around which assignments for

- a first course in computer system could be designed,
- (c) its shortcomings and omissions (and their possible solutions), particularly as support for advanced systems courses,
- (d) the problems encountered in the implementation of a working system.

The preceding chapters have largely ignored the feasibility of such a system from an economics point of view. No really detailed analysis can be made comparing the cost of using the manufacturer's software with that of using SUPERSALT, for running assignments, until the extended SALT system has been used in this capacity. However, the author has noted that the use of OS-MVT software for assignments in a previously presented systems course cost between \$1.00 and \$10.00 per run, depending on the assignment. They also required the use of temporary disk files. The far more extensive SUPERSALT assignments outlined in Chapter IV would, in the author's opinion, cost between \$.10 and \$2.00 per run and would require I/O only on system input and output devices. While it is not the function of this thesis to establish the limits of economic feasibility for any course, it may be concluded that SUPERSALT offers a more economic and more versatile mechanism for running assignments for a systems course than has hitherto been employed.

The author further concludes that this investigation has established both the feasibility of creating such a system and the desirability of using it as the foundation on which a course in computer systems could be structured.


BIBLIOGRAPHY

1. Achtemichuk, L. (1967), "Survey of Computer Science Courses at Canadian Institutes", Supplement to the Quarterly Bulletin, vol. 7, 4.
2. ACM Curriculum Committee on Computer Science, (1968), "Curriculum 1968 (Recommendations for Academic Programs in Computer Science)", Communications of the ACM, vol. 11, 3, pp. 151-197.
3. BCS Working Party 7, (1967), "Survey of Computer Science Courses", Computer Bulletin, vol. 11, 1, pp. 31-42.
4. Center for Computer and Information Sciences, (1968), SOS: The Brown University Student Operating System, Brown University, Providence, Rhode Island.
5. Dutton, J.B. (1971), "Programming the SUPERSALT Computer", Computing Center Publication, University of Alberta, (in preparation).
6. Dutton, J.B. (1970), "A Pseudo Clock for SALT", Computing Review, University of Alberta Student Chapter of the Association for Computing Machinery, vol. 3, pp. 33-46.
7. Dutton, J.B. (1969), "Student Assembler Language Translator System Description", Computing Center Publication, University of Alberta.

8. Dutton, J.B. (1971), "SUPERSALT System Description", Computing Center Publication, University of Alberta (in preparation).
9. Dutton, J.B. and Penny, J.P. (1970), "Undergraduate Education in the Fundamental Concepts of Large Computer Systems", Proceedings IFIP World Conference on Computer Education, Amsterdam, vol. 2, pp. 185-189.
10. Easton, L.S. and Penny, J.P. (1969), "Student Assembler Language Reference Manual", Department of Computing Science, University of Alberta.
11. Jones, P.D. (1967), "Operating System Structures", Proceedings IFIP Congress, Edinburgh, pp. 29-38.
12. Schorr, H. and Waite, W.M. (1967), "An Efficient Machine-Independent Procedure for Garbage Collection in Various List Structures", Communications of the ACM, vol. 10,8, pp. 501-506.
13. Wile, D.S., Munchk, R.G. and Van Dam, A. (1967), "The Brown University Student Operating System", Proceedings of the ACM National Meeting, pp. 427-439.

APPENDIX I

ADDITIONS TO THE MACHINE INSTRUCTION SET



The execution of this instruction signals the addressed channel (if available) to fetch eight bytes from the address given in the Channel Address Word and to use these eight bytes as a Channel Command Word directed to the addressed device.

The D1, B1 fields of the instruction are not used to generate a memory address. Instead, the low order sixteen bits of the sum formed by the addition of the contents of register B1 and the contents of the D1 field identify the channel which is to fetch a CCW and the device to which the CCW is to apply.

The I/O operation is initiated if:


- (a) the addressed channel exists and is not currently active performing a previously initiated I/O operation, and
- (b) the addressed device exists and is not currently active performing a previously initiated I/O operation.

A channel and device are considered busy until the interrupt marking the end of the operation has occurred. Thus, if that interrupt is masked off, the channel and device are unavailable even though the I/O operation may have been completed at the device.

The resulting condition code indicates the status of the I/O attempt.

Condition Code:

- 0 - I/O operation initiated successfully and channel now in operation
- 1 - the addressed channel found to be busy controlling I/O on the addressed device
- 2 - the addressed channel found to be busy performing I/O on some device other than the addressed device
- 3 - the I/O could not be initiated - a Channel Status Word indicating the reason is stored

Program Interrupt:**Privileged Operation**

The execution of this instruction gives indication of the state of the addressed channel and device by causing the condition code to be set and the instantaneous Channel Status Word to be stored.

The D1(B1) fields of the instruction are not used to generate a memory address. Instead, the low order sixteen bits of the sum formed by the addition of the contents of register B1 and the contents of the D1 field identify the channel and device that are to be tested.


Condition Code:

- 0 - both the addressed channel and device presently free
- 1 - the addressed channel found to be busy controlling I/O on the addressed device

- 2 - the addressed channel found to be busy performing I/O on some device other than the addressed device
- 3 - no test could be made but a Channel Status Word is stored anyway, indicating the reason

Program Interrupt:

Privileged Operation



The execution of this instruction causes the immediate termination of any I/O operation in progress on the addressed channel and device, and causes the instantaneous Channel Status Word to be stored. The results of the instruction execution are indicated by the setting of the condition code.

The D1(B1) fields of the instruction are not used to generate a memory address. Instead, the low order sixteen bits of the sum formed by the addition of the contents of register B1 and the contents of the D1 field identify the channel and device to which the instruction applies.

Condition Code:

- 0 - the addressed device found to be not engaged in I/O activity
- 1 - the I/O is complete, but the device and channel are not free since a masked I/O interrupt has been stacked in hardware before execution of the HIO - the interrupt can only be cleared by setting the system mask in the current PSW to unmask the addressed channel

2 - the I/O operation in progress on the addressed device has been halted and that device and the addressed channel are now free

3 - an invalid condition exists and is specified in detail in the stored Channel Status Word

Program Interrupt:

Privileged Operation




The byte at the memory location designated by the operand address replaces the system mask of the current PSW.

Condition Code: Unchanged

Program Interrupts:

Privileged Operation

Addressing



The eight bytes at the full-word aligned location specified by the operand address completely replaces the current PSW which is lost. The sixteen general purpose registers are loaded from a memory location reserved for this purpose (see Appendix II). Instruction execution continues under control of the newly loaded PSW at the location contained in the instruction address field of the new PSW. This address is not checked for validity during execution of the LPSW but is checked during execution of the following instruction. The LPSW

is executed under control of the old current PSW which must have the CPU in the supervisor state. The newly loaded PSW, however, can place the CPU into the problem state without contradicting the privileged nature of the LPSW instruction.

Condition Code:

The condition code is set according to bits 34 and 35 of the newly loaded PSW.

Program Interrupts:

Privileged Operation

Addressing

Specification

Note: In the event of a program interrupt during execution of an LPSW instruction, the PSW stored in the Program Old PSW slot is the old current PSW.

APPENDIX II

SUPERSALT HARDWARE CHARACTERISTICS

1. Specification of I/O Device Characteristics

Hardware specification for the devices and channels attached to the SUPERSALT CPU are detailed on the \$SALT card in special option fields of the form:

$$\text{DEVICE} = \left\{ \begin{array}{l} \text{READER} \\ \text{PRINTER} \end{array} \right\} , \text{addr}, \text{datarate}, \text{avedelay}$$

where

READER - indicates an input unit record device of fixed record length of 80 bytes

PRINTER - indicates an output unit record device of fixed record length of 120 bytes plus one byte carriage control

addr - a three character hexadecimal constant of the form abb where "bb" is the device address, "a" is a number 0 to 7 and is the address of the channel to which the device is connected

datarate - the rate of transfer of information through the device and down the channel in bytes per second

avedelay - the average delay time in microseconds before the command is carried out at the device

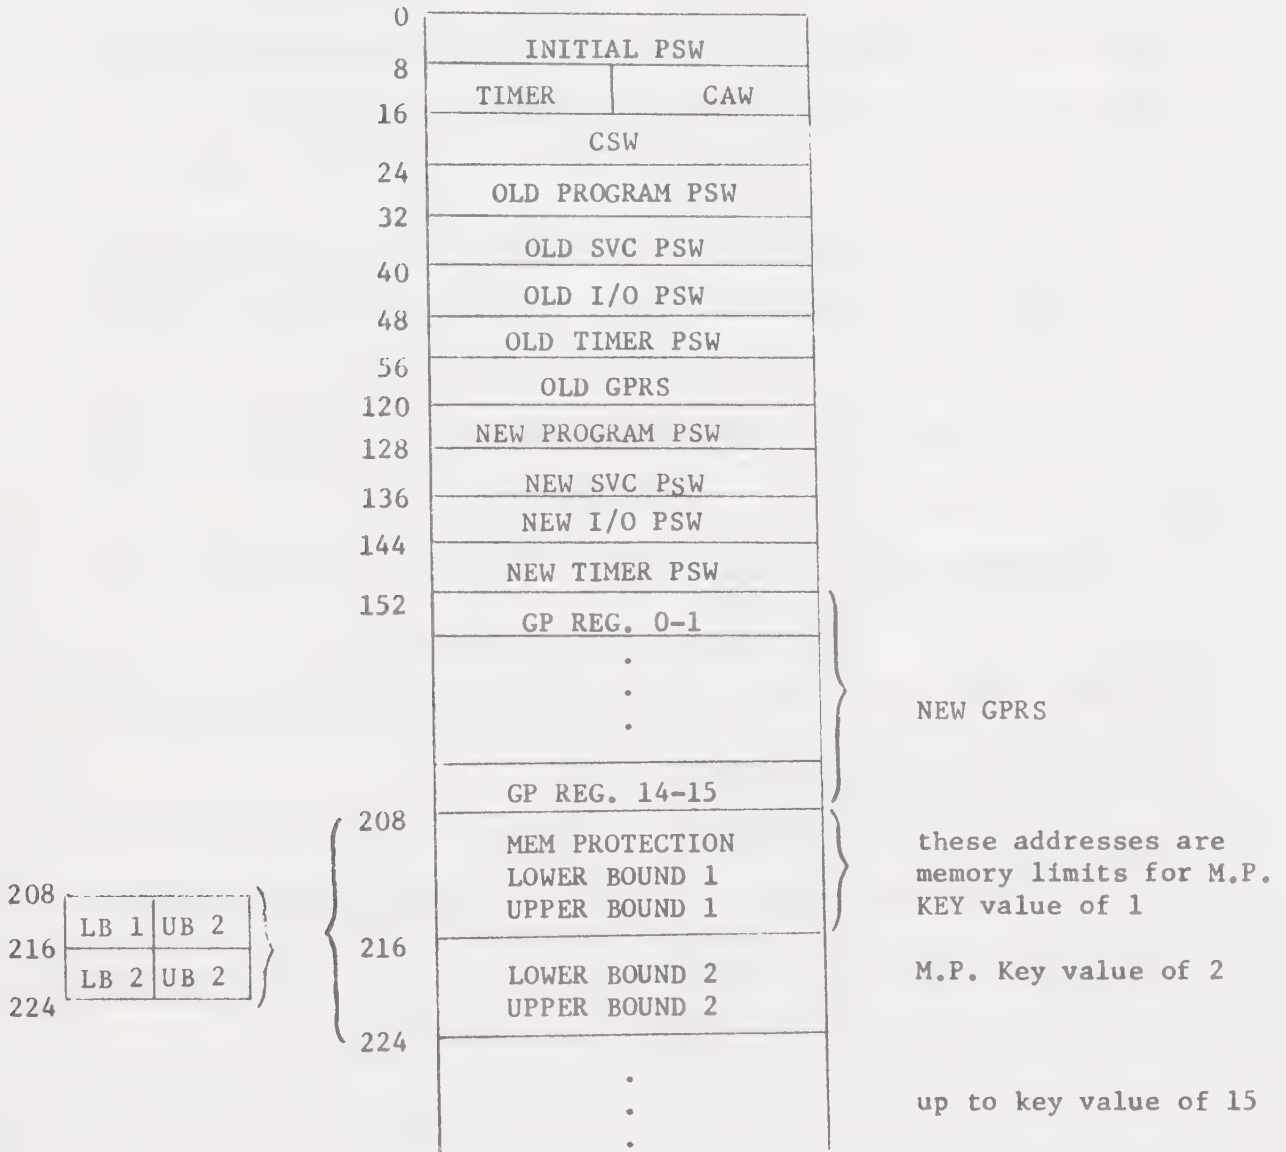
2. Specification of CPU Reserved Memory Format

The user may provide his own PSW's by including on the \$SALT card the following option field:

PSW=addr

where

addr - is the relative address of a memory area in the user's SALT program which is considered to have the following format:



where

INITIAL PSW - is the PSW loaded at the beginning of execution of the SALT program.

TIMER - is a full word that may be set to any value by the programmer. It is decremented in synchronous with execution of SALT instructions and causes a timer interrupt whenever the value goes from zero to a negative value or when it goes from the maximum negative number to zero (not currently implemented).

CAW - (Channel Address Word) set by the programmer to the address of a Channel Command Word before an SIO instruction is executed.

CSW - Channel Status Word; see No. 3 following.

OLD 'X' PSW's - the location where the current PSW is stored during an 'X' class interrupt.

NEW 'X' PSW's - the location where a new PSW is loaded from for an 'X' class interrupt. See No. 4 below for PSW format.

G.P. REG 0-15 - contents of user's general purpose registers; stored by hardware during an interrupt.


MEM PROTECTION LOWER BOUND 'X' - the address below which memory is store-protected when the PSW Memory Protect Key value is 'X'.

MEM PROTECTION UPPER BOUND 'X' - the address above which memory is store-protected when the PSW Memory Protect Key value is 'X'.

- 39 - the status information just stored is the instantaneous status of a continuing I/O operation
- 40 - unused
- 41 - unused
- 42 - the addressed channel does not exist
- 43 - the addressed device is not one of those controlled by the addressed channel
- 44 - command not recognized by the addressed device
- 45 - 'control' command contains invalid modifier bit combination
- 46 - command not recognized by the addressed channel
- 47 - attempt to read past physical end of device

Residual Byte Count - the difference between the count value obtained from the count field of the current CCW and the number of bytes transferred down the channel (up to the time the CSW was stored) under control of that CCW's command.

4. Program Status Word Format



A one bit in any of the following positions indicates:

bit(s)

0 - any I/O interrupt associated with an I/O operation on this

channel is postponed until a new PSW in which this bit is zero is loaded

1 - 7 same for channels 1 to 7

14 - the CPU is in the WAIT state (RUN state when the bit is zero)

15 - the CPU is in the PROBLEM state (SUPERVISOR state when the bit is zero)

16 - 31 the code for the interrupt that caused the storing of this PSW, that is; for an SVC interrupt the SVC code, for an I/O interrupt the channel/device address, and for a program interrupt the program interrupt code

32 - 33 instruction length code for the current instruction

00 - 2 bytes	RR
01 - 4 bytes	RX
10 - 4 bytes	RS-SI
11 - 6 bytes	SS

34 - 35 condition code

40 - 63 address of current instruction

APPENDIX III

THE CHANNEL COMMAND WORD

1. The Assembler Instruction

name CCW command,addr,flag,count

where

COMMAND is an absolute expression specifying the command to be executed by the channel. The channel command set consists of

0000 0001 - write
 0000 0010 - read
 0000 1000 - transfer in channel
 0000 MM11 - control printer

MM = 00 skip page

01 space 3

10 space 2

11 space 1

ADDR is a relocatable expression specifying the address where:

- (a) the channel is to transmit data to (READ),
- (b) the channel is to retrieve data from (WRITE),
- (c) the channel is to retrieve the next CCW from (TRANSFER IN CHANNEL).

FLAG is an absolute expression that specifies the flags for bits 32-36.


and zeros for bits 37-39 of the machine command (see below).

COUNT is an absolute expression that specifies the number of bytes of

data to be acted upon by this command.

2. The Machine Command Format

The SALT assembler creates from the CCW statement an eight-byte field, aligned on a full-word boundary, having the following format:



where a one bit in the following position indicates:

bit

- 32 - Chain Data: this command is to apply (without reference to the CPU) to the COUNT and DATA ADDRESS fields of the next contiguous CCW, following completion of the current CCW.
- 33 - Chain Command: the channel is to fetch and execute (without reference to the CPU) the next contiguous CCW, following completion of the current CCW. (However, if the current CCW has caused EOF on an input device, an I/O interrupt terminates automatic channel fetching even though command chaining is specified.)
- 34 - Unused
- 35 - Skip: for this CCW, data transmission across the channel is to be suppressed; however, the command is to be carried out at the device.
- 36-39 - 0 (unused)

NOTE: Only command chaining can be carried across a TRANSFER IN CHANNEL.

APPENDIX IV

ADDITIONS TO THE SYSTEM MACRO SET

1. Execute Channel Program

name	EXCP	ccbname
+	L	1,=A(ccbname)
+	SVC	6

Description: The EXCP macro requests the SALT monitor (or a user-supplied supervisor if the "PSW=" option was specified on the \$SALT card) to initiate an I/O operation on the issuing program's behalf. The Channel Control Block whose address is placed in register 1 is assumed to contain the

- (a) address of the channel to be used,
- (b) address of the required device on that channel, and
- (c) the address of the CCW to be used by the channel.

2. Test and Wait on I/O

name	WAIT	ccbname
+	L	1,=A(ccbname)
+	TM	8(1),X'01'
+	BZ	*+6
+	SVC	7

Description: The WAIT macro is used to impose whatever synchronizing constraints are required between program instruction execution

and channel operation. The Channel Control Block, whose address is placed in register 1, is assumed to contain a status bit set to one by the monitor (or supervisor) whenever the CCB was named in an EXCP for an I/O operation that is currently in progress and set to zero following the I/O interrupt terminating that I/O operation. When the WAIT macro is issued for a CCB flagged as in use, the WAIT SVC is executed, providing the monitor (supervisor) with the opportunity to wait on completion of the I/O.

3. Channel Command Block

name	CCB	X'deviceadr',ccwname	
+	DC	A(ccwname)	
+	DC	X'deviceadr'	
+	DC	2X'00'	RESIDUAL COUNT
+	DC	2X'00'	STATUS
+	DC	2X'00'	SPARE

where

(a) deviceadr is four hexadecimal digits of the form ccdd

cc - channel address 00 to 07

dd - device address 00 to FF

(b) ccwname is the symbolic address of the CCW associated with this CCB.

Description: The CCB provides a means of communicating shared information between the issuing program and the resident monitor or supervisor. A given CCB is specific to a given unit (deviceadr) and a

given CCW chain (ccwname); however, more than one CCB can refer to the same unit or CCW chain. The residual count field is used by the monitor or supervisor, on completion of the I/O operation referencing this CCB as a field in which to store the difference between the count value obtained from the count field of the last CCW in the chain and the number of bytes transmitted down the channel during execution of that CCW. The status field is used by the monitor or supervisor to reflect to the issuing program the status of the last completed I/O operation referencing this CCB, if the unit is presently free, or the current I/O operation, if it is busy. A one bit in the following status field positions indicate:

Bit

- 0 - 1 (unused)
- 2 - the current CCW specifies chain data
- 3 - the current CCW specifies chain command
- 4 - the last completed I/O operation caused end of file
- 5 - the count value obtained from the count field of the final CCW in the chain of CCW's last completed was exhausted before the end of the physical block
- 6 - the last I/O operation caused the end of the physical block at the device before the count value, obtained from the final CCW in the chain of CCW's last completed, was exhausted
- 7 - this CCB is being used by an I/O operation currently in progress.

APPENDIX V

A SAMPLE SALT PROGRAM

The following computer listing is that of a SALT program run on the SUPERSALT machine. The program contains an elementary supervisor consisting of simple interrupt handlers to service the needs of a problem program employing a variety of different channel programs. The hardware configuration was specified on the \$SALT card with fields specifying:

DEVICE=READER,1C1,120000,100

DEVICE=PRINTER,2C2,220000,200

PSW=4

LJC	OBJECT CODE	ADDR1	ADDR2	STMT	SOURCE STATEMENT
0000F4	D203 300C 30F0 0000C	0000C	000F0	48	*** TIMER INTERRUPT HANDLING ROUTINE
0000FA	P200 3034 00034			49	TIMER(4),HIVALU BOOST TIMER
0000FO	3B0ACA00			50	LPSW OTPSW CONTINUE WITH PREVIOUS PSW
				51	HIVALU DC F'10C0000000'
				52	*** PROGRAM INTERRUPT HANDLING ROUTINE
				53	PROGPROC HI0 X'02C2' DUMP OLD PSW'S AND GPHS
0000F4	9E00 02C2		C02C2	54	DUMP
0000F8	4110 3C00		00300	55	LA 1,SUPFR
0000FC	4100 0100		00100	56	LA 0,256
000100	0A05			57	SVC 5
000102	18A4			58	SR R10,R10
000104	42A0 301F		0001F	59	IC R10,OPPSW*3 GET INT CODE
000108	89A3 0002		00002	60	SLL R10,2 TIMES 4
00010C	58A4 310C		0010C	61	L R10,MADTAB(R10) MESSAGE ADDR
000110	D213 3101 A000	00161	00000	62	MVC MSGSLDT(20),0(R10)
000116	9E00 02C2		002C2	63	HI0 X'02C2' SCRAP PREVIOUS IO
00011A	9E00 01C1		001C1	64	HI0 X'01C1'
00011E	41A0 3144		00144	65	LA R10,PINTCCW
000122	5CA0 3010		00010	66	ST R10,CAW
000126	9C00 02C2		002C2	67	BC 7,OUT+8
00012A	4770 314C		0014C	68	MVC NIOPSW+4(4),OUT+4
00012E	D203 3090 313C	00090	0013C	69	OUT LPSW
000134	8200 3138		00138	70	DS OF
000138	FF020000			71	DC X'FFC20000' SUP/WAIT STATES-INT'S ENABLED-PROT KEY 0
00013C	C000014C			72	DC A(*+4) EOJ ADDR
				73	EIJ
000140	0A00			74	SVC 0
000144	CF07015400000000			75	SPACE1,ERLINE,CCOM,0
00014C	0100015400000078			76	WRITE,ERLINE,NDFLAGS,120
000154	C1C2C5D5C40C04			77	C'REND DUE TO ' 107C' '
000161	4C40404040404040			78	INT CODE 1
				79	INT CODE 2
				80	INT CODE 3
				81	INT CODE 4
				82	INT CODE 5
				83	INT CODE 6
				84	INT CODE 7
				85	INT CODE 8
0001CC	00000000			86	C'DIVIDE EXCEPTION
0001D0	000001FC			87	C*PRIVILEGED OPERATION
0001D4	00002004			88	C*INVALID OP CODE
0001D8	00000218			89	C*OVERFLOW
0001DC	0000022C			90	C*ADDRESS ALIGNMENT
0001E0	00000240			91	C*REG SPECIFICATION
0001E4	00000254			92	C*OPERAND ALIGNMENT
0001E8	00000268			93	C*PROTECT VIOLATION
0001FC	0000027C			94	C*INVALID IO CONDITION
				95	
0001F0	C4C9E5C9C4C540			96	DIVMSG DC
000204	D709C9F5C9C3C5			97	PRIVMSG DC
000218	C9D5E5C1D3C9C4			98	OPERMMSG DC
00022C	D6E5C5D9C6D3D6			99	OVFLMSG DC
000240	C1C4C400C5E2F2			100	SPECMMSG DC
000254	D9C5C740E207C5			101	REGMSG DC
000268	D6D7C5D0C135C4				UPRNDMSG DC
00027C	D7D0AE3C5C3E3				PROTMMSG DC
000290	C9D5E5C1D3C9C4				10MSG DC

000244	1RAA				102	*** SUPERVISOR CALL (SVC) INTERRUPT HANDLING ROUTINE
000246	43A0	3027			103	SVCTBL SR R10,R10
000248	89A0	0002			104	IC R10,OSVCPSPW+3 GET SVC NO.
00024E	59A0	320C			105	SLL R10,2 TIMES 4
000252	4720	3482			106	C R10,SVCTBL TEST VALIDITY
000256	52FA	328C			107	RH R10,SVCTBL
00025A	C7FF				108	RAVSV
000260	00003E0				109	L R15,SVCTBL(R10)
000262	00003E2				110	BR R15
000264	0000320				111	SVCTBL DC AISVCEQJ SVC 0
000268	0000398				112	DC AISVCREAD 1
00026C	00003A4				113	DC AISVCRWITE 2
000270	00003BC				114	DC AISVCGTI 3
000274	00003CE				115	DC AISVCTC 4
000278	00003D0				116	DC AISVCDUMP 5
00027C	00003E0				117	DC AISVCEXCP 6
000280	00003E2				118	DC AISVCGWITE 7
000284	00003E4				119	SVCTBL DC AISVCTHLE-SVCTBL-4)
000288	00003E6				120	***
000292	00003E8				121	* END OF JOB-REFLECT SVC TO SALT MONITOR
000296	00003EA				122	SVCTBL ENJ
000300	00003EC				123	+SVCTBL SVC 0
000304	00003EE				124	***
000308	00003F0				125	* READ SVC - I/O NOT OVERLAPPED WITH PROBLEM PROG INST EXECUTION
000312	00003F2				126	SVCTBL NI SVCEXCP,X'00',DISALLOW EXCP
000316	00003F4				127	MVC R0R0(W+1(3),0GPR5+5 READ INTO ADDR FROM GPR1
000320	00003F6				128	LA R10,PRACCM
000324	00003F8				129	ST R10,CAN
000328	00003FA				130	SIN X'01C1
000332	00003FC				131	BC 7,PRACCM
000336	00003FE				132	IORTN MVC NIOPSN(8),OSVCPSPW
000340	0000400				133	MVC NIOPSN(64),0GPR5 RETURN TO USER AFTER INT
000344	0000402				134	GI OSVCPSPW+1,X'02' WITH USER'S REGISTERS
000348	0000404				135	LPSW OSVCPSPW PROH/WAIT STATES-INT'S ENABLED
000352	0000406				136	CCW READ,R0VCCW,NOFLAGS,80 WAIT ON I/O
000356	0000408				137	READER LA R10,IMWSS
000360	000040A				138	CANJDR
000364	000040C				139	***
000368	000040E				140	* WRITE SVC - I/O NOT OVERLAPPED WITH PROGRAM INSTRUCTION EXEC
000372	0000410				141	SVCTBL NI SVCEXCP,X'00',DISALLOW EXCP
000376	0000412				142	LA R10,PRACCM
000380	0000414				143	ST R10,CAN
000384	0000416				144	L R10,0GPR5+4
000388	0000418				145	CLI 0(R10),C'0' TEST USER CARRIAGE CONTROL
000392	000041A				146	RE SPACES1
000396	000041C				147	CLI 0(R10),C'0'
000400	000041E				148	RE SPACES2
000404	0000420				149	CLI 0(R10),C'0'
000408	0000422				150	BE SPACES3
000412	0000424				151	CLI 0(R10),C'1'
000416	0000426				152	BE SKIPGE
000420	0000428				153	SPACES1 MVI PTRCCW,X'0F'
000424	000042A				154	B DDIO
000428	000042C				155	SPACES2 MVI PTRCCW,X'0B'

SOURCE STATEMENT

STMT

ADDR1

ADDR2

OBJECT CODE

LDC

00035C	47F0	336C			156	SPACES3	B	DOIO	
000360	9207	3388		0036C	157	MVI	B	PTIRCCW,X'07*	
000364	47F0	336C		00388	158	B	DOIO		
000368	9203	3388		0038C	159	MVI	PTIRCCW,X'03*		
00036C	41A4	0021		00021	160	LA	R10,1(R10)	USER BUFFER + 1	
000370	59A0	3070		00030	161	ST	R10,SUPER		
000374	0202	3391	3001	00001	162	MVC	PTIRCCW+9(3),SUPER+1		
00037A	90C0	0202		002C2	163	SIO	X'02C2*		
00037E	47F0	3318		00318	164	BC	7,WRITEERR		
000382	47F0	32FC		002FC	165	R	10RTN		
000388	000018840	000000			166	CCW	X'03*,PTIRCCW,CCOM,C	CONTROL CCW	
000390	0100038840	00000078			167	CCW	WRITE,PTIRCCW,NDFLAGS,120		
000319					168	WRITEERR	EOU	READERR	
					169	*			
00039A	5800	303C		0003C	170	* CHAR TO INTEGER CONVERSION SVC			
00039C	5810	3040		00040	171	SVCCTI	L	R0,0GRPS	
0003A0	0A02				172	L	R1,0GRPS+4		
0003A2	5020	3044		00044	173	SVC	3	CALL ON SALT MONITOR FOR CONVERSION	
0003A6	8200	3024		00024	174	ST	R2,0GRPS+8	REFLECT RESULT TO USER	
					175	LPSW	DSVCPSW		
					176	*			
0003AA	5800	303C		0003C	177	* ITC CONVERSION SVC			
0003AE	5810	3040		00040	178	SVCITC	L	R0,0GRPS	
0003B2	0A04				179	L	R1,0GRPS+4		
0003B4	5020	3044		00044	180	SVC	4	CALL SALT MONITOR FOR CONVERSION	
0003B8	8200	3024		00024	181	ST	R2,0GRPS+8		
					182	LPSW	DSVCPSW		
					183	*			
0003BC	9E00	02C2		002C2	184	* DUMP SVC			
0003C0	5500	303C		0003C	185	HIO	X'02C2*	READY PRINTER	
0003C4	5810	3040		00040	186	L	R0,0GRPS		
0003C8	CA05				187	L	R1,0GRPS+4		
0003CA	8200	3024		00024	188	SVC	5	CALL SALT MONITOR FOR DUMP	
					189	LPSW	DSVCPSW		
					190	*			
0003CE	45E0	346C		0046C	191	* EXECUTE CHANNEL PROGRAM SVC			
0003D2	9101	100A		0000A	192	SVCXCP	BAL	R14,GETUNIT	
0003D6	4710	3472		00472	193	TM	B(P1),X'01*	CCR BUSY?	
0003DA	9000	AC00		00000	194	BU	BSYCCR	YES	
0003DE	4780	343A		0043A	195	TIO	R(RIC)		
0003E2	4700	341F		0041F	196	BC	R,SNEWIO		
0003E6	41A0	345E		0045E	197	BC	6,0NEWIO		
0003FA	47F0	311C		00110	198	LA	R10,RADIOMSG	FATAL I/O CONDITION	
0003FE	C9D5EC10	3C9C4			199	R	CANJIB		
000402	41A0	340A		0040A	200	RADIOMSG	DC	C,INVALID IO CONDITION*	
000406	47F0	3110		00110	201	BSYCCR	LA	R10,BSYCCRM	
00041A	C09A40E3	09E840			202	R	CANJIB		
00041E	41A0	3426		00426	203	BSYCCRM	DC	C,IO TRY ON BUSY CCB *	
000422	47F0	3110		00110	204	QNEWIO	LA	R10,BSYUNITM	COULD QUEUE IO AT THIS POINT
000426	C9D540E3	09E840			205	R	CANJIB		
00043A	5821	C00C		0000C	206	BSYUNITM	DC	C,IO TRY ON BUSY UNIT *	
00043E	5020	3010		00010	207	SNEWIO	L	R2,0(R1)	GET CCB ADDR
					208	ST	R2,CAN	POINT CAN AT CCW	
000442	9601	1008		00008	209	OI	B(R1),X'01*	MARK CCB AS BUSY	

LCC	OBJECT CODE	ADDP1	ADDR-2	STMT	SOURCE STATEMENT
000446	5940 3520		00520	210	C R10,P10B
000447	4770 3456		00456	211	R1E *+12
000448	5210 3524		00524	212	ST R1,R10B+4
000449	4780 345A		0045A	213	B *+8
000450	5940 3520		00520	214	ST R1,P10B+4
000451	4780 3450		00450	215	STU O(R10)
000452	4780 3458		00458	216	BC B,*+10
000453	0201 1008	3018	00018	217	MVC B(2,P1),CSW+4
000454	0200 3024		00024	218	LPSW OSVCPSW
000455	5310 3040		00040	219	* GETUNIT L R1,UGPRS+4
000456	1922 3000		00000	220	SR R2,P2
000457	5200 3000		00000	221	ST R2,SUPER
000458	0201 3002	1004	00002	222	MVC SUPER+2(2),4(R1)
000459	5300 3000		00000	223	L R10,SUPER
000460	07FE			224	PR R14
000461				225	
000462	4500 3460		00460	226	* WAIT SVC
000463	5940 3460		00000	227	SVCWAIT R14,GETUNIT
000464	4740 3492		00492	228	TIO O(R10)
000465	8200 3024		00024	229	BC 4,*+9
000466	5940 3520		00520	230	LPSW OSVCPSW
000467	4770 3442		00442	231	C R10,R10B
000468	3524 3446		00524	232	PNE *+12
000469	4780 3446		00446	233	OI R10H+4,X*80
000470	3520 3520		00520	234	B *+8
000471	4780 3446		00446	235	OI P10H+4,X*80
000472	94FF 1008		00008	236	NI R(1),X*FE
000473	9602 3025		00025	237	OSVCPSW+1,X*C2
000474	8200 3024		00024	238	UI OSVCPSW+1,X*C2
000475	4140 348A		0048A	239	LPSW OSVCPSW
000476	4780 3110		00110	240	BADSVC LA R10,MSVC
000477	005F5C103C9C4			241	B CANJOB
000478				242	MSVC DC C*INVALID SVC CODE
000479				243	*** I/O INTERRUPT HANDLING ROUTINE
000480	1922 3000		00000	244	IOPROC SR R2,R2
000481	5020 3002	302E	00002	245	ST R2,SUPER
000482	0201 3000		00000	246	MVC SUPER+2(2),OIOPSW+2
000483	5940 3520		00520	247	L R10,SUPER
000484	4770 34FF		004FF	248	C R10,R10B
000485	4120 3520		00520	249	RNE *+12
000486	4780 34F2		004F2	250	LA R2,R10B
000487	4120 3528		00528	251	R *+8
000488	0201 1006	301A	00006	252	LA R2,P10B
000489	0201 1008	3018	00008	253	L R1,4(R2)
000490	0201 3020		00020	254	MVC 6(2,R1),CSW+6
000491	4710 350E		0050E	255	MVC 8(2,R1),CSW+4
000492	8200 3020		00020	256	TM OIOPSW+1,X*02
000493	8200 302C		0002C	257	BO *+8
000494	1211 350A		0050A	258	YES
000495	4720 302D		0002D	259	OIOPSW
000496	94FD 302D		0002D	260	LPSW
000497				261	LTR
000498				262	BP
000499				263	IOUT
000500					NI
000501					OIOPSW+1,X*FD
000502					TURN WAIT BIT OFF

MINI SUPERVISOR-INTERRUPT HANDLING ROUTINES						PAGE 6
VERSION 2 U OF A SALT 4 DECEMBER 1970						
LCC	OBJECT CODE	ADDR1	ADDR2	STMT	SOURCE STATEMENT	
000518	9400 2004		00004	264	NI 4(P2),X'00'	REMOVE FLAG FROM I/O BLOCK
00051C	8200 302C		0002C	265	LPSW 010PSW	RETURN TO PROBLEM PROG
				266	*	
000520				267	R10B	
000520	CCCC01C1			268	DS OF	
000524	CCCC0000			269	DC X'CCCC001C1'	READER ADDR
000528				270	DS A(0)	CURRENT CCR ADDR
000528	000002C2			271	DS OF	
00052C	CCCC0000			272	DC X'000002C2'	PRINTER ADDR
000530				273	DC A(0)	CURRENT CCR ADDR
000530				274	USENTRY 3	
				275	DS OF	
				276	*** END OF SUPERVISOR	

LOC	URJECT CODE	ADDR1	ADDR2	STMT	SOURCE STATEMENT
000530	05A0			279	BALR B10,C
000532				279	USING *,B1C
				280	*
				281	* SIMPLE CCW READ FOLLOWED BY SKIP PAGE CNTRLCCW CHAINED
				282	* TO A SIMPLF CCW WRITE
				283	EXCP RDRCCB1
000532	5810 A5F2		00P24	284 +	L 1,=A(RDRCCB1)
000536	CA06			285 +	SVC 6
				286	WAIT RDRCCB1
000538	5810 A5F2		00R24	287 +	L 1,=A(RDRCCB1)
000540	9131 1008		00C08	288 +	TM 8(1),X*31*
000542	4780 AC14		00546	289 +	BZ **6
000544	CA07			290 +	SVC 7
000546	9108		006CC	291	TM RDRCCB1*8,X*08*
000548	4710 A15E		00690	292	ENDPRUG
00054E	D24F A3E6	A396 00918	00P08	293	MVC LINE(90),CARD
				294	EXCP PTRCCB1
000554	5810 A5F6		00R28	295 +	L 1,=A(PTRCCB1)
000558	CA06			296 +	SVC 6
				297	WAIT PTRCCB1
00055A	5810 A5F6		00E28	298 +	L 1,=A(PTRCCB1)
00055E	9131 1008		000C9	299 +	TM 8(1),X*31*
000562	4780 AC36		00568	300 +	BZ **6
000566	CA07			301 +	SVC 7
				302	*
				303	* UNCHAINED PRINTER CONTROL CCW FOLLOWED BY A SIMPLE WRITE CCW
000568	D24F A3E6	A53A 00918	00A6C	304	MVC LINE(90),MSG1
00056E	5810 A5FA		0082C	305	EXCP PTRCCB2
000572	0A06			306 +	L 1,=A(PTRCCB2)
				307 +	SVC 6
				308	WAIT PTRCCB2
000574	5810 A5FA		00R2C	309 +	L 1,=A(PTRCCB2)
000578	9131 1008		00008	310 +	TM 8(1),X*31*
00057C	4780 A050		00582	311 +	BZ **6
000580	CA07			312 +	SVC 7
				313	EXCP PTRCCB3
000582	5810 A5FE		00B30	314 +	L 1,=A(PTRCCB3)
000586	CA06			315 +	SVC 6
				316	WAIT PTRCCB3
000588	5810 A5FE		00R30	317 +	L 1,=A(PTRCCB3)
00059C	9131 1008		00C08	318 +	TM 8(1),X*31*
00059D	4780 AC64		00596	319 +	BZ **6
000594	CA07			320 +	SVC 7
				321	*
				322	* SIMPLE WRITE WITH NO PRINTER CNTRL CAUSING PRINTER OVERSCORE
000596	D24F A3E6	A588 00918	00A8A	323	MVC LINE(80),MSG2
00059C	5810 A5FE		00B30	324	EXCP PTRCCB3
0005A0	CA06			325 +	L 1,=A(PTRCCB3)
				326 +	SVC 6
				327	WAIT PTRCCB3
0005A2	5810 A5FE		00R30	328 +	L 1,=A(PTRCCB3)
0005A6	9131 1008		00C08	329 +	TM 8(1),X*31*
0005AA	4780 A07E		00580	330 +	BZ **6
0005AE	CA07			331 +	SVC 7


```

332 *-----
333 * CHAIN DATA IN THE FORM OF A SCATTER READ FOLLOWED BY A SPACE 3
334 * CONTROL CCW COMMAND CHAINED TO A SIMPLE WRITE CCW
335 * MVI LINE,C ,
336 MVC LINE+1(119),LINE BLANK LINE
337 EXCP RORCCB2
338 + L 1,A(RORCCB2)
339 + SVC 6
340 WAIT RORCCB2
341 + L 1,A(RORCCB2)
342 + TM 8(1),X*31,
343 + RZ **5
344 + SVC 7
345 TM RORCCB2+8,X*08, EDF?
346 BD ENOPROG YES
347 EXCP PTRCCB5
348 + L 1,A(PTRCCB5)
349 + SVC 6
350 WAIT PTRCCB5
351 + L 1,A(PTRCCB5)
352 + TM 8(1),X*31,
353 + RZ **5
354 + SVC 7
355 *-----
356 * A READ CHANNEL PROGRAM USING THE SKIP DATA OPTION FOLLOWED BY A
357 * WRITE CHANNEL PROGRAM USING SKIP DATA
358 * MVI CARD,C ,
359 MVC CARD+1(79),CARD BLANK INPUT BUFFER
360 MVC CARD*20(20),C***NOT OVERWRITTEN***
361 EXCP RORCCB3
362 + L 1,A(RORCCB3)
363 + SVC 6
364 WAIT RORCCB3
365 + L 1,A(RORCCB3)
366 + TM 8(1),X*31,
367 + RZ **6
368 + SVC 7
369 DUMP CARD*80
370 + LA 1,CARD
371 + LA 0,80
372 + SVC 5
373 MVI LINE,C,X,
374 MVC LINE+1(119),LINE FILL LINE WITH X'S
375 EXCP PTRCCB6
376 + L 1,A(PTRCCB6)
377 + SVC 6
378 WAIT PTRCCB6
379 + L 1,A(PTRCCB6)
380 + TM 8(1),X*31,
381 + RZ **6
382 + SVC 7
383 *-----
384 * CPU/CHANNEL OVERLAP DEMONSTRATED BY INITIATING CHANNEL TO
385 * WRITE FROM A BUFFER FILLED WITH A'S WHILE CPU INSTRUCTION

```


LDC	SUBJECT CODE	ADDR1	ADDR2	STMT	SOURCE STATEMENT
000636	5910 A626			384 *	EXECUTION PROGRESSIVELY FILLS BUFFER FROM HIGH END WITH B'S
000637	5910 A626			387	EXCP PTRCCB4
000638	CA3A		00B58	389 +	L 1,=A(PTRCCB4)
000639	4150 C077		00G77	389 +	SVC 6
000640	4140 A45E		00G90	390	LA R5,119
000641	1A45			391	DESTROY LA R4,LINEA
000642	1A45			392	AR R4,R5
000643	1200 4000 A62A	03000	00B5C	393	MVC 0(1),R4),=C*B*
000644	4650 A10E		00G40	394	BCT R5,DESTROY
000645	5910 A626		00B5R	395	WAIT PTRCCB4
000646	9131 1008		00C08	396 +	L 1,=A(PTRCCB4)
000647	4780 A12C		00G5E	397 +	TM 8(1),X*31*
000648	0A07			399 +	BZ **6
000649				399 +	SVC 7
000650				400 *	-----
000651				401 *	SOPHISTICATED CHANNEL PROGRAM EMPLOYING TRANSFER IN CHANNEL
000652				402 *	CONCURRENT WITH A READ THAT READS NO DATA BUT CAUSES END OF FILE
000653	5910 A62F		00B60	403	EXCP PTRCCB7
000654	CA3A			404 +	L 1,=A(PTRCCB7)
000655	5910 A62F		00B60	405 +	SVC 6
000656	5910 A62F		00B60	406	EXCP RD*CCB1
000657	5910 A62F		00B60	407 +	L 1,=A(RD*CCB1)
000658	5910 A62F		00B60	408 +	SVC 6
000659	5910 A62F		00B60	409	WAIT RD*CCB1
000660	5910 A62F		00B60	410 +	L 1,=A(RD*CCB1)
000661	5910 A62F		00B60	411 +	TM 8(1),X*31*
000662	5910 A62F		00B60	412 +	BZ **6
000663	5910 A62F		00B60	413 +	SVC 7
000664	5910 A62F		00B60	414	WAIT PTRCCB7
000665	5910 A62F		00B60	415 +	L 1,=A(PTRCCB7)
000666	5910 A62F		00B60	416 +	TM 8(1),X*31*
000667	5910 A62F		00B60	417 +	BZ **6
000668	5910 A62F		00B60	418 +	SVC 7
000669	5910 A62F		00B60	419	TM RD*CCB1*8,X*08*
000670	5910 A62F		00B60	420	EDF? YES
000671	5910 A62F		00B60	421	ENDPRG NO-CAUSE SALT INTERRUPT
000672	5910 A62F		00B60	422	DC X*0000*
000673	5910 A62F		00B60	423	ENDPRG DUMP
000674	5910 A62F		00B60	424	EDFMSG,ENDCCBS-EUFMSG
000675	5910 A62F		00B60	425 +	LA 1,EDFMSG
000676	5910 A62F		00B60	426 +	LA 0,ENDCCBS-EUFMSG
000677	5910 A62F		00B60	427	SVC 5
000678	5910 A62F		00B60	428 +	FOJ
000679	5910 A62F		00B60	429	SVC 0
000680	5910 A62F		00B60	430 *	-----
000681	5910 A62F		00B60	431 *	CHANNEL CONTROL BLOCKS WITH DUMP FORMAT MESSAGES
000682	5910 A62F		00B60	432 *	READER
000683	5910 A62F		00B60	433	DS
000684	5910 A62F		00B60	434	ENFMSG DC C**CONTENTS OF CCB S AT RDR EDF**
000685	5910 A62F		00B60	435	DC C**RDRCCB1*
000686	5910 A62F		00B60	436	RDRCCB1 CCB X*01C1*,RDRCCW1
000687	5910 A62F		00B60	437	+RDRCCB1 DC A(RDRCCW1)
000688	5910 A62F		00B60	438	SIMPLE READ
000689	5910 A62F		00B60	439	
000690	5910 A62F		00B60	440	
000691	5910 A62F		00B60	441	
000692	5910 A62F		00B60	442	
000693	5910 A62F		00B60	443	
000694	5910 A62F		00B60	444	
000695	5910 A62F		00B60	445	
000696	5910 A62F		00B60	446	
000697	5910 A62F		00B60	447	
000698	5910 A62F		00B60	448	
000699	5910 A62F		00B60	449	
000700	5910 A62F		00B60	450	

LDC	OBJECT CODE	ADDR1	ADDR2	STMT	SOURCE STATEMENT	
0006C8	01C1			438 +	DC X'01C1'	
0006CA	0000			439 +	DC 2X'00'	
0006CC	0000			440 +	DC 2X'00'	
0006CE	40404040404040			441	DC 14C'	
0006D0	00000000000000			442	DC C'DRDRCCB2'	SCATTER READ
0006E4	00000000			443	R'DRCCB2 CCB X'01C1',RDRCCW2	
0006E8	01C1			444 +	R'DRCCB2 DC A(P'DRCCW2)	
0006EA	0000			445 +	DC X'01C1'	
0006EC	0000			446 +	DC 2X'00'	
0006EE	40404040404040			447 +	DC 2X'00'	
0006F0	00000000000000			448	DC 14C'	
000704	00000000			449	DC C'DRDRCCB3'	
000708	01C1			450	R'DRCCB3 CCB X'01C1',RDRCCW3	SKIP DATA READ
00070A	0000			451 +	R'DRCCB3 DC A(P'DRCCW3)	
00070C	0000			452 +	DC X'01C1'	
00070E	40404040404040			453 +	DC 2X'00'	
000710	00000000000000			454 +	DC 2X'00'	
000712	00000000000000			455	DC 14C'	
000714	00000000			456 *	PRINTER	
000716	00000000000000			457	DC C'PTRCCB1'	
000718	00000000			458	PTRCCB1 CCB X'02C2',PTRCCW1	SKIP PAGE WITH SIMPLE WRITE
00071A	0000			459 +	PTRCCB1 DC A(P'PTRCCW1)	
00071C	0000			460 +	DC X'02C2'	
00071E	0000			461 +	DC 2X'00'	
000720	0000			462 +	DC 2X'00'	
000722	40404040404040			463	DC 14C'	
000724	00000000000000			464	DC C'PTRCCB2'	
000726	00000000			465	PTRCCB2 CCB X'02C2',PTRCCW2	SINGLE SPACE 2 CONTROL CCW
000728	0000			466 +	PTRCCB2 DC A(P'PTRCCW2)	
00072A	0000			467 +	DC X'02C2'	
00072C	0000			468 +	DC 2X'00'	
00072E	40404040404040			469 +	DC 2X'00'	
000730	00000000000000			470	DC 14C'	
000732	00000000000000			471	DC C'PTRCCB3'	
000734	00000000			472	PTRCCB3 CCB X'02C2',PTRCCW3	SINGLE WRITE CCW
000736	0000			473 +	PTRCCB3 DC A(P'PTRCCW3)	
000738	0000			474 +	DC X'02C2'	
00073A	0000			475 +	DC 2X'00'	
00073C	40404040404040			476 +	DC 2X'00'	
00073E	00000000000000			477	DC 14C'	
000740	00000000000000			478	DC C'PTRCCB4'	
000742	0000			479	PTRCCB4 CCB X'02C2',PTRCCW4	SPACE 1 CNTRL WITH SIMPLE WRITE
000744	0000			480 +	PTRCCB4 DC A(P'PTRCCW4)	
000746	0000			481 +	DC X'02C2'	
000748	0000			482 +	DC 2X'00'	
00074A	40404040404040			483 +	DC 2X'00'	
00074C	00000000000000			484	DC 14C'	
00074E	00000000000000			485	DC C'PTRCCB5'	
000750	00000000			486	PTRCCB5 CCB X'02C2',PTRCCW5	SPACE 3 CNTRL WITH SIMPLE WRITE
000752	0000			487 +	PTRCCB5 DC A(P'PTRCCW5)	
000754	0000			488 +	DC X'02C2'	
000756	0000			489 +	DC 2X'00'	
000758	40404040404040			490 +	DC 2X'00'	
00075A	00000000000000			491	DC 14C'	

PROBLEM PROGRAM - TRY VARIOUS CCW S					PAGE 12	
SOURCE STATEMENT					VERSION 2 U OF A SALT 4 DECEMBER 1970	
LOC	OBJECT CODE	ADDR1	ADDR2	STAT	SOURCE STATEMENT	
000001				542	* CCW COMMANDS	
000002				543	WRITE EQU 1	X'01'
000003				544	READ EQU 2	X'02'
000004				545	TIC EQU 8	X'08'
000005				546	SPACE1 EQU 15	X'0F'
000006				547	SPACE2 EQU 11	X'0B'
000007				548	SPACE3 EQU 7	X'07'
000008				549	SKIPAGE EQU 3	X'03'
000009				550	* CCW FLAGS	
000010				551	NOFLAGS EQU 0	
000011				552	SKIP EQU 16	X'10'
000012				553	CCOM EQU 64	X'40'
000013				554	CDAT EQU 128	X'80'
000014				555	CDATSK EQU 144	X'90'
000015				556	* BUFFERS	
000016				557	CAPO DC 800'	
000017				558	LINE DC 1200'	
000018				559	LINEA DC 1200'A'	
000019				560	BLANKS DC 1000'	
000020				561	* MESSAGES AND MISC	
000021				562	MSG1 DC C*THIS LINE PRINTED BY UNCHAINED WRITE CCW PRECEDED BY *	
000022				563	MSG2 DC C*UNCHAINED SPACE 2 CONTROL*	
000023				564	DC 200*X'	
000024				565	DC C* THIS LINE PRINTED WITH NO PREVIOUS PRINTER CONTROL*	
000025				566	DC 350'	
000026				567	R0 EQU 0	
000027				568	R1 EQU 1	
000028				569	R2 EQU 2	
000029				570	R3 EQU 3	
000030				571	R4 EQU 4	
000031				572	R5 EQU 5	
000032				573	R10 EQU 10	
000033				574	R10 EQU 10	
000034				575	R14 EQU 14	
000035				576	R15 EQU 15	
000036				577	*-----	
000037				578	END	
000038				579	=====	
000039				580	=====	
000040				581	=====	
000041				582	=====	
000042				583	=====	
000043				584	=====	
000044				585	=====	
000045				586	=====	
000046				587	=====	
000047				588	=====	
000048				589	=====	
000049				590	=====	

0 STATEMENT(S) FLAGGED IN THIS ASSEMBLY

B29972